NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA

THESIS

INTERNETWORKING: INTEGRATING IP/ATM LAN/WAN SECURITY

by

Ronald M. Dennis

September, 1996

Thesis Advisor: Don Brutzman
Associate Advisor: Rex Buddenberg

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
September 1996

3. REPORT TYPE AND DATES COVERED
Master's Thesis

4. TITLE AND SUBTITLE
INTERNETWORKING: INTEGRATING IP/ATM LAN/WAN SECURITY

5. FUNDING NUMBERS

6. AUTHOR(S)
Ronald M. Dennis

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT

Computer and network security is a complex problem that is not solely restricted to classified computer systems and networks. Accelerating trends in networking and the emphasis on open interoperable networks has left many unclassified systems vulnerable to a wide variety of attacks. Computer and network professionals must understand the scope of security, recognize the need for security even in unclassified systems, and then take appropriate action to protect their systems.

Transmission of static passwords in plaintext over the Internet is one of the most widely publicized network vulnerabilities. One-time password mechanisms (such as S-Key) or other secure network access mechanisms (such as Kerberos) have been recommended to improve access security for computer systems connected to the Internet.

This thesis examines many of the issues that must be addressed when assessing the need for computer and network security. This work provides the results of a site security survey for the unclassified IP/ATM LAN in the Systems Technology Lab (STL) at the Naval Postgraduate School (NPS). These results highlight new security vulnerabilities and strengths that occur when standard Internet Protocol (IP) local-area networks (LANs) are internetworked with Asynchronous Transfer Mode (ATM) wide-area networks (WANs). Finally, we examine the feasibility of using the Kerberos authentication protocol for remote plaintext password protection and provide recommendations for additional work.

14. SUBJECT TERMS network security, computer security, Kerberos, Asynchronous Transfer Mode (ATM), internetworking, risk management.

15. NUMBER OF PAGES 191

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT
Unclassified

18. SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

19. SECURITY CLASSIFICATION OF ABSTRACT
Unclassified

20. LIMITATION OF ABSTRACT
UL

Approved for public release; distribution is unlimited.

INTERNETWORKING: INTEGRATING IP/ATM LAN/WAN SECURITY

Ronald M. Dennis
Lieutenant, United States Navy
B.S., University of Washington, 1988

Submitted in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
September 1996

ABSTRACT

Computer and network security is a complex problem and one that is not solely restricted to classified computer systems and networks. Accelerating trends in networking and the emphasis on open and interoperable networks has left many unclassified systems vulnerable to a wide variety of attacks. Computer and network professionals must understand the scope of security, recognize the need for security for even unclassified systems and then take steps to protect their systems.

Transmission of static passwords in plaintext over the Internet is one of the most widely publicized network vulnerabilities. One-time password mechanisms (such as S-Key) or other secure network access mechanisms (such as Kerberos) have been recommended to improve access security for computer systems connected to the Internet.

This thesis examines many of the issues that must be addressed when assessing the need for computer and network security. This work provides the results of a site security survey for the unclassified IP/ATM LAN in the Systems Technology Lab at the Naval Postgraduate School. These results highlight new security vulnerabilities and strengths that occur when standard Internet Protocol (IP) local-area networks (LANs) are internetworked with Asynchronous Transfer Mode (ATM) wide-area networks (WANs). Finally, we examine the feasibility of the Kerberos authentication protocol for remote plaintext password protection and provide recommendations for additional work.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

AAL	ATM Adaptation Layer
ACL	Access Control List
ADP	Automated Data Processing
AIS	Automated Information System
ARPANet	Advanced Research Projects Agency Network
ATM	Asynchronous Transfer Mode
AUP	Acceptable Use Policy
CAT	Common Authentication Technology
CERT	Computer Emergency Response Team
CGI	Common Gateway Interface
CNO	Chief of Naval Operations
CNS	Cygnus Network Security
DAC	Discretionary Access Controls
DES	Digital Encryption Standard
DMS	Defense Message System
DNS	Domain Name System
DoD	Department of Defense
DoN	Department of the Navy
DOS	Disk Operating System
DSS	Digital Signature Standard

E-911	Emergency 911
FIPS	Federal Information Processing Standards
ftp	file transfer protocol
GBS	Global Broadcast System
IAB	Internet Architecture Board
ID	Identification
IETF	Internet Engineering Task Force
IIRG	Information Infrastructure Research Group
IP	Internet Protocol
IPSEC	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IS	Information System
ISO	International Standards Organization
ISP	Internet Service Provider
I-WAY	Information Wide-Area Year
KDC	Kerberos Key Distribution Center
LAN	Local-Area Network
MAC	Mandatory Access Controls
MBone	Multicast Backbone
MD5	Message Digest 5
MIT	Massachusetts Institute of Technology

NAVPGSCOLINST	Naval Postgraduate School Instruction
NCSC	National Computer Security Center
NIST	National Institute of Standards and Technology
NPS	Naval Postgraduate School
NSFNET	National Science Foundation Network
OAM&P	Operations, Administration, Maintenance & Provisioning
ODU	Old Dominion University
OMB	Office of Management and Budget
OSI	Open Systems Interconnection
OTP	One-Time Password
PGP	Pretty Good Privacy
pop	post office protocol
PSN	Public Switched Network
PVC	Permanent Virtual Circuit
QoS	Quality of Service
rcp	remote copy
RFC	Request for Comments
rlogin	remote login
rm	remove (files/directories)
RSA	Rivest, Shamir and Adleman (cryptographic algorithm)
rshell	remote shell

SECNAVINST	Secretary of the Navy Instruction
SEI	Software Engineering Institute
set-uid	set user identification
SGI	Silicon Graphics Incorporated
SHA	Secure Hash Algorithm
SONET	Synchronous Optical Network
SPAWAR	Space and Electronic Warfare Center
STL	Systems Technology Lab
STP	Signal Transfer Point
SVC	Switched Virtual Circuit
TCB	Trusted Computing Base
TCP	Transport Control Protocol
TCSEC	Trusted Computer Security Evaluation Criteria
TDN	Tactical Data Network
telnet	terminal emulation application
tftp	trivial file transfer protocol
TGS	Ticket Granting Server
TGT	Ticket Granting Ticket
UDP	User Datagram Protocol
UPS	Uninteruptible Power Supply
VC	Virtual Circuit
WAN	Wide-Area Network

WB	Whiteboard
WWW	World Wide Web

ACKNOWLEDGEMENTS

This has been a long and arduous journey, often with no clear path to follow. Without the support of a great many people I could not have finished. I wish to thank Professor Don Brutzman for his persistent enthusiasm, support and motivation in the face of setbacks and disappointments. I also wish to thank Rex Buddenberg for his candid opinions, essential insights and often off-the-wall “sea stories” throughout the thesis process.

Don McGregor and the rest of the STL staff deserve more credit than I can possibly give in this short space. Without Don’s help installing and configuring the software, there would be no results (positive or negative). Thanks also to Professor Glen Wheless at the Center for Coastal Physical Oceanography at Old Dominion University in Norfolk Virginia for his willingness and patience with installing and testing the software.

I express my appreciation for Kathy Powers and the people at Cygnus Support for their help and flexibility in customizing a technical support package to meet our needs on such short notice. Thanks also to Terry Williams (the STL Network Manager) for coming up with the funding and to Alexandra Sumners with the Undersea Warfare Academic Group administrative support staff for expediting our purchase request.

There are so many other people to thank and so little space in which to do it but I want to thank my family and my close friends (you know who you are) for their never ending love, support and encouragement. Lastly I want to thank God for seeing me through this endeavor, for His unfailing love and grace, and for the gifts that He has given me.

I. INTRODUCTION

A. PROBLEM DEFINITION

Computer and network security is a complex problem and one that is no longer unique to classified networks. The Internet and internetworked computer systems in general offer incredible financial, educational, scientific and recreational opportunities. While the ease with which individuals and organizations are able to transmit and retrieves data and information around the world is certainly attractive, it does not come without pitfalls. Threats to computer system security and network security abound. The problem faced by many organizations, including the Department of Defense (DoD), is being able to take advantage of this global internetworking trend without compromising proprietary or sensitive information, or falling prey to malicious attacks against the normal operation of their networked computer systems.

Personal data, credit card numbers, proprietary business information and sensitive research data must be protected both while stored in host systems and while being transmitted over the network. Many people believe that new networking technologies such as Asynchronous Transfer Mode (ATM) will significantly reduce the threats to networked computer system security. However, current networking technology will remain in place for the foreseeable future. Regardless of whether the physical technology and protocols used for the Internet backbone are based on routed IP packets or switched ATM circuits, internetworking trends will continue. We are beginning to see internetworked hybrids of both approaches, along with a need for a keener focus on computer and network security.

Computer and network security is primarily a management issue, one that historically has been a low priority. Managers must understand internetworking trends and the threats to networked computer system security. They must recognize the need to secure even unclassified networks. This thesis focuses on the security threats and vulnerabilities associated with internetworking a Internet Protocol (IP) local-area network (LAN) and an Asynchronous Transfer Mode (ATM) wide-area network (WAN) at the Naval Postgraduate School (NPS).

B. MOTIVATION

1. Growth of the Internet

The incredible growth of the Internet and the technological advances in modern computer network communications has made global communication and information interchange a reality. Individual and organizational computer systems are accessible to millions of other computers and individuals where they were once relatively sheltered. Over the past 25 years the Internet has experienced incredible growth. Table 1[1] shows that in the last year alone, the number of computer hosts connected to the Internet has nearly doubled and the number of networks, as indicated by the Network Class figures, has *more than doubled* (Lottor, 1996). These statistics should be considered minimums due to incomplete data collection as discussed in (Lottor, 1992).

[1] Data in Tables 1 and 2 were produced by Network Wizards and is used by permission. The original survey data are available at *http://www.nw.com/*

Date	Hosts	Domains	Network Class		
			A	B	C
Jan 96	9,472,000	240,000	92	5655	87,924
Jul 95	6,642,000	120,000	91	5390	56,057
Jan 95	4,852,000	71,000	91	4979	34,340

Table 1. Number of hosts, domains and networks advertised in the Domain Name System (DNS). After (Lottor, 1996).

While the Internet was initially implemented in the 1970's as a demonstration project using new networking technology for the U.S. Government (Rowe, 1995; Krol, 1994; Claffy, *et al.*, 1994), Table 2 shows that the Internet now connects millions of computers on tens of thousands of networks belonging to governments, organizations, businesses and academia from around the world (Lottor, 1996). This growth will undoubtedly continue.

Domain	Hosts
Commercial (*com*)	2,430,954
Educational (*edu*)	1,793,491
Networks (*net*)	758,597
Germany (*de*)	452,997
United Kingdom (*uk*)	451,750
Canada (*ca*)	372,891
Government (*gov*)	312,330
Australia (*au*)	309,562
Organizations (*org*)	265,327
U.S. Military (*mil*)	258,791

Total:	7,406,690

Table 2. Host distribution by top-level domain name (Lottor, 1996).

DoD use of the Internet also continues to increase. The U.S. DoD (i.e. the *mil* domain) ranks tenth with regard to the number of hosts connected to the Internet, and the Pentagon continues working to take advantage of internetworking capabilities and emerging technologies. All of the military services emphasize the need for integrating computer and communication networks and for using cutting edge technology (Sullivan, 1995; CNO, 1995; Fogleman and Widnall, 1995). These integrated networks (or at least the information they transport) will need to be secure to be effectively used by the military commander in the field (Nierle, 1996).

2. The Need for Computer Security

The need for computer and network security is obvious when speaking about classified defense information or when viewed from a national defense perspective. However, although NPS is a military installation, its mission is not to provide national defense but to enhance it by providing graduate and professional education for U.S. and allied military officers (NPS, 1995). If NPS only has an indirect effect on national security, why should it (or any other educational institution) be concerned with computer system and network security? Why should business and industry (aside from those defense contractors developing classified technology) be concerned with networked computer system security? These questions deserve analysis.

The original "Internet," which consisted initially of the DoD's Advanced Research Project Agency Network (ARPANet) and later the National Science Foundation Network of supercomputers (NSFNET), was developed primarily to allow collaborating researchers to exchange files and electronic mail messages among one another. To this day, the

Internet is essential for collaborative research among academia, industry and the military, but the Internet user community is no longer composed strictly of benevolent academicians and researchers exchanging data and communications for the good of science. Banking services, electronic commerce, international business transactions and a host of commercial on-line services all compete for network bandwidth. (Krol, 1994; Boucher, 1994) Not all Internet users are model citizens; some seek to exploit system and network weaknesses either out of malice or else for personal or professional gain.

These malicious users are not necessarily external to one's own organization. Statistics show that external attacks account for less than 3% of reported losses while malicious insiders account for approximately 19% and mistakes made by well-intentioned insiders account for approximately 65% of reported losses (Wong and Watt, 1990). Therefore the networked computer system security umbrella must protect an organization against threats from inside as well as outside the organization. The important point is that even if only 0.01% of the estimated 20 million Internet users (Blum, 1995-1996; Boucher, 1994) is attempting to gain unauthorized access or otherwise cause damage to your computer systems/networks or the data they contain, whether they are inside or outside your organization, there are still 20,000 wolves at the door. The most frightening aspect of these wolves is that they can be at more than one networked door at a time, they can use a variety of techniques to gain access to computer systems, and their actions can have long-lasting effects.

If the mere possibility of a computer security incident doesn't convince computer system users, administrators and managers of the need for computer security, the applicable laws and regulations should. Table 3 lists a number of laws and directives that are relevant to this thesis. These laws and directive proscribe various security requirements for federal computer systems. The following bullets briefly describe each of these laws and directives.

- The Privacy Act of 1974 was enacted largely in response to growing public concern over the privacy of personal information gathering in various government databases in the early 1970's (Russell and Gangemi, 1991).

- Office of Management and Budget (OMB) Circular A-71 requires federal agencies under the jurisdiction of the executive branch to establish and maintain computer security programs (Russell and Gangemi, 1991; DoN, 1991).

- OMB Circular A-130 (Management of Federal Information Resources, 1985) extended the A-71 requirement to all federal agencies.

- The Computer Fraud and Abuse Act of 1986 (Public Law 99-474; amendment to title 18 U.S. Code, section 1030) was enacted to combat the growing number of computer crimes in the early 1980's — crimes that cost federally insured financial institutions an estimated $730 million in 1985. This act also prohibits unauthorized access to computer systems owned or operated by or on behalf of the Federal Government. (Toensing, 1986).

- The Computer Security Act of 1987 (Public Law 100-235) requires all U.S. government computer systems containing sensitive information to have a customized security plan (Russell and Gangemi, 1991; DoN, 1991).

- The Privacy Act of 1974
- OMB Circular A-71
- OMB Circular A-130
- The Computer Fraud and Abuse Act of 1986
- The Computer Security Act of 1987
- DoD Directive 5200.28
- SECNAVINST 5239.2

Table 3. Relevant laws and directives governing computer security.

- Department of Defense (DoD) Directive 5200.28 "Security Requirements for Automated Information Systems (AISs)" outlines the security requirements for all automated information systems operated by or for the DoD (SPAWAR, 1993; DoN, 1991).

- Secretary of the Navy Instruction (SECNAVINST) 5239.2 "Department of the Navy Automated Information Systems" provides specific guidelines that apply to automated information systems operated by or for the Department of the Navy (SPAWAR, 1993; DoN, 1991).

C. THESIS OBJECTIVES

There are two primary objectives for this thesis. The first objective is to document the results of a site security survey for the unclassified IP/ATM LAN/WAN at NPS and identify areas that are deficient. The second objective is to provide a secure remote login capability to the IP/ATM LAN/WAN at NPS for mobile users and for research partners at Old Dominion University in Norfolk Virginia.

D. THESIS STRUCTURE

1. Scope

An overwhelming variety of subjects might be examined when studying networked computer security. We have focused on practical and experimentally verifiable issues of immediate interest. The scope of this thesis is limited to a discussion of the security issues relevant to the integrated IP/ATM LAN/WAN at NPS, documenting the results of a site security survey for the IP/ATM LAN/WAN at NPS, and implementing and testing the Kerberos authentication and authorization protocol to improve remote user access security of the IP/ATM LAN/WAN at NPS.

2. Organization

This introductory chapter defines the general problem and explains the motivation for the research described and documented in this thesis. Chapter II presents background information necessary to understand the remainder of this work and describes some related internetworking and networked computer system security research. Chapter III presents a detailed problem statement. Chapter IV discusses many of the issues that managers must address when determining the type and extent of controls required to adequately secure their networked computer system. Chapter V describes the Kerberos authentication and authorization protocol. Chapter VI documents the results of a site security survey for the IP/ATM LAN/WAN at NPS. Chapter VII describes the steps taken to implement the Kerberos software in the NPS Systems Technology Lab (STL). Research conclusions and recommendations for further work are presented in Chapter VIII.

II. BACKGROUND AND RELATED WORK

A. INTRODUCTION

This chapter provides background information and briefly discusses related internetworking and computer security research. Section B defines the concepts and purpose of computer security for our context. Section C discusses the International Standards Organization (ISO) Open Systems Interconnection (OSI) logical reference model for network communications and how computer security can be implemented within the framework of the OSI model. The Internet protocol model is briefly discussed and contrasted with the OSI model in Section D. A general discussion of ATM, its relationship to the Internet model and its vulnerabilities with regard to security is provided in Section E. The chapter concludes with a discussion of related internetworking and security research in Section F.

B. COMPUTER SECURITY BACKGROUND

Computer security is a broad subject. Before we can discuss the subject we must first define precisely what it is we are discussing. Subsection 1 defines computer security and explains its purpose. Subsection 2 describes the Department of the Navy's computer security program (DoN, 1991). Subsection 3 clarifies the terms threat, vulnerability, risk and control as used in the context of networked computer security.

1. Definition and Purpose

Ask ten different computer professionals for the definition and purpose of computer security and you are likely to receive ten different variations on a theme. The following quotations serve to illustrate this point.

> "[Computer security is] the protection of the computer resources against accidental or intentional disclosure of confidential data, unlawful modification of data or programs, the destruction of data, software or hardware, and the denial of one's own computer facilities . . .." (Palmer and Potter, 1989, pg. 12)
>
> "[Computer security is] the protection of . . . systems and their data, against any accidental or deliberate compromise of proprietary data and programs . . . and assuring their continuous availability to users . . . through built-in system resilience features to minimise short term disruptions, and contingency planning, to mitigate protracted system stoppage from natural or man-made disasters." (Wong and Watt, 1990, pg. 23)
>
> Computer security is the protection of your "computer and everything associated with it . . .." (Russell and Gangemi, 1991)

Even the military services do not agree precisely on what computer security is. The U.S. Air Force defines computer security as the protection of information confidentiality, integrity and availability (Fleming, 1993), while the U.S. Navy defines automated information system (AIS) security as the policies and procedures taken to "protect AISs against unauthorized (accidental or intentional) destruction, modification, disclosure and denial of service" (DoN, 1991, para. 1.1).

Despite subtle variations in the definitions given above, each is an appropriate definition of computer security. Nevertheless some concepts are more appropriate than others for today's internetworked computing environment. Where once computer security focused on securing access to a central mainframe computer, security in today's distributed computing environment requires a much broader scope. One cannot consider the security of an individual computer without considering the security of the information that is processed by that computer regardless of whether the information is stored on the computer itself or on a secondary storage device. Likewise in the case of networked computers, one cannot consider computer and information security without considering the security of the network over which the information is transmitted. (Cobb, 1992; Baskerville, 1988; Reeves, 1983)

This thesis focuses on a local-area network (LAN) owned and operated by the Naval Postgraduate School (NPS) and therefore uses the computer security definition outlined in the Department of the Navy (DoN) Automated Information Systems (AIS) Security Guidelines (DoN, 1991). The definition applies to the hardware, the software and the data, regardless of whether such data is in storage, is being processed by the central processing unit or is being transmitted over a network. Thus the terms automated information system (AIS) security, system/network security and internetworked computer security are used interchangeably throughout the remainder of this thesis.

2. DoN AIS Security Guidelines

The Department of the Navy divides its overall security program into seven interdependent elements. These elements are shown in Table 4. Although each element is described and governed by a separate instruction, all categories overlap to a great extent. For example, the DoN AIS Security Guidelines (the first element in Table 4) document the roles, responsibilities and requirements for ensuring adequate information security, network and data communications security, personnel security, physical security and

- AIS Security
- Information Security
- Communications Security
    - Cryptographic security
    - Physical security
    - Transmission security
    - Emission security
- Personnel Security
- Physical Security
- Emanations Security (TEMPEST)
- Network Security
    - Network integrity
    - Data authentication
    - Field integrity
    - Non-repudiation
    - Protocol-based protection
    - Connection confidentiality
    - Access control

Table 4. DoN security program elements (DoN, 1991, para. 1.5).

emanations security as each applies to automated information systems (AISs). Although it may seem that some of these elements are meaningless outside the context of AISs, each individual element does in fact encompass a much broader scope.

The primary purpose of the DoN AIS Security Guidelines is to assist Navy commands and "activities in meeting the *intent* of the DoN AIS Security policies" (emphasis added) (DoN, 1991, para. 1.4). Some readers may argue that the list in Table 4 is incomplete in that it does not include system availability. This is indeed an important factor of computer security as discussed in Chapter IV. However, although the list does not explicitly include it, system availability is implicitly considered in the risk analysis methodology outlined in the guidelines.

The DoN AIS Security Guidelines are not merely an implementation checklist. They are a management tool commanding officers or other responsible individuals can use to determine the required level of security for the activity's AISs. The guidelines provide a methodology and sample documentation that commands and activities can use, both to analyze the security of their AISs and to determine if adequate threat controls are in place to operate those AISs at an acceptable level of risk. What constitutes the appropriate amount of security and the acceptable level of risk are management decisions that will vary from system to system.

3. Threats, Vulnerabilities, Risks and Controls

Whichever definition is chosen, AIS security is the art of mitigating threats by eliminating or reducing system vulnerabilities through the application of controls in order to minimize threat impact and the risk of loss. Unfortunately these three terms — threat, vulnerability and risk — are used often indiscriminately. For example, the DoN AIS Security Guidelines treat the terms "threat" and "vulnerability" synonymously. However, the three terms are clearly distinct. The following paragraphs define these terms.

Threats are possible pitfalls or dangers confronting an information system, no matter how likely or unlikely they may be. For example, fire is a threat to overall information system security. In accordance with our definition of AIS security, fire can result in damaged or destroyed hardware/software/data or denial of system services. (Fitzgerald, 1993; Russell and Gangemi, 1991) As with fire, threats need not be directed at or against the information system *per se*, but may simply be a product of the environment that could cause coincidental damage to the system. A threat corresponds to a system vulnerability only in the absence of adequate proactive controls.

A *vulnerability* represents a security weakness in the system that might result in a potential loss if that weakness is exploited by a threat (Fitzgerald, 1993; Fleming, 1993; Russell and Gangemi, 1991; Wong and Watt, 1990; Palmer and Potter, 1989). The frequency at which a threat may occur is not a factor when determining system vulnerabilities. For example, if an information system does not have an operative or adequate fire suppression system or other fire protection mechanism, this lack of fire protection is a system vulnerability no matter how rarely a fire may actually occur. (Wong and Watt, 1990)

Risk is a measure of the possible impact or loss that may result if a threat actually exploits a system vulnerability. The level of risk will depend on the controls in place, the value of the assets affected and the likelihood that the threat will actually occur. (Cobb, 1992; Wong and Watt, 1990; Palmer and Potter, 1989; DoN, 1991)

Controls are the countermeasures, protection mechanisms or safeguards implemented to counter threats to a system and to reduce the risk of loss. Traditionally controls have one of three purposes: prevention, detection or correction. Preventative controls are designed to prevent or hinder a threat from occurring. Detective controls are designed to alert users and/or system administrators if a particular threat or event occurs. Corrective controls are designed to assist users and/or system administrators in recovering after a threat or event occurs. (Fitzgerald, 1993; Russell and Gangemi, 1991; Wong and Watt, 1990; Baskerville, 1988)

Now that computer security has been defined, its purpose explained and some frequently misunderstood terms clarified, a brief discussion of network communication protocols is in order. The next three sections briefly discuss the ISO OSI reference model, the Internet protocol (TCP/IP) suite and emerging Asynchronous Transfer Mode (ATM) technology.

C. ISO OSI REFERENCE MODEL

It seems almost canonical that any discussion of networking must include a discussion of the International Standards Organization (ISO) Open Systems Interconnection (OSI) logical reference model for network and data communications. While it is important to have a basic understanding of the functionality associated with each layer in the OSI model, one must realize that the model is only conceptual. It is not an implementation guide. The OSI model merely provides a convenient way for researchers, technicians and computer professionals to discuss network functionality. Subsection 1 describes the functionality of each layer of the OSI model and Subsection 2 describes the security services associated with each layer of the model.

1. Network Functionality in the OSI Model

The ISO OSI logical reference model is shown in Figure 1. Each layer of the OSI model outlines a particular function with regard to data and network communications.

The bottom three layers — the physical layer, the data link layer and the network layer — work together to actually access the network and get the data onto the "wires." The *physical layer* is responsible for bit encoding (i.e. transforming binary code into the appropriate form for transmission on the physical medium). This layer is responsible for

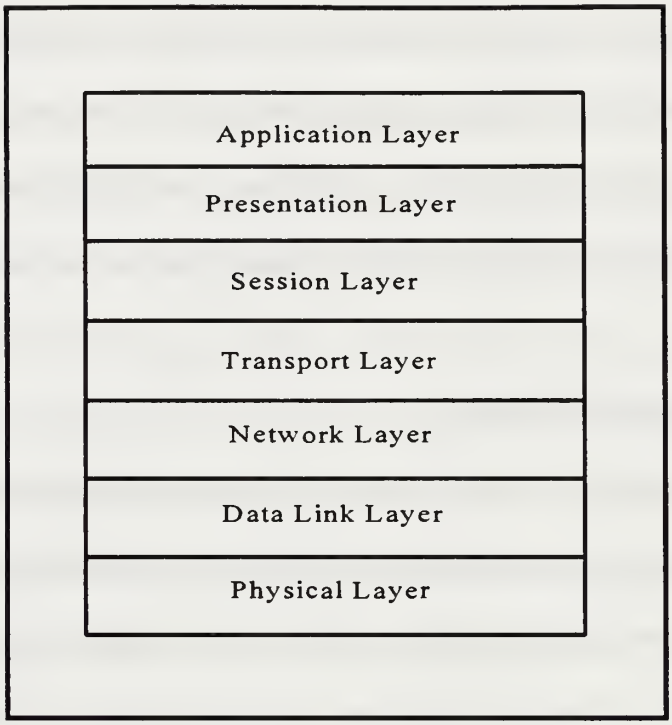

Figure 1. ISO OSI reference model.

making the physical, electrical and/or mechanical connection to the physical medium and is concerned with transmitting and receiving the data stream onto and from the network. The *data link layer* is responsible for point-to-point message synchronization and integrity and is concerned with encapsulating the data into data units and calculating error checksums. The *network layer* is responsible for efficient routing across a switched network and is concerned with logical-to-physical address resolution. (Fitzgerald, 1993; Russell and Gangemi, 1991; Wong and Watt, 1990)

The top four layers — the transport layer, the session layer, the presentation layer and the application layer — work together to establish a connection and to format the data for presentation to the lower three layers. The *transport layer* is responsible for end-to-end data integrity, message assembly/disassembly and connection management. The *session layer* is responsible for establishing and maintaining the connection for a particular communication session. The *presentation layer* is responsible for formatting data for user applications and other devices on the network. Finally, the *application layer* is responsible for application software (e.g. ftp, word processing etc.) and the end user's interface to the network. (Fitzgerald, 1993; Russell and Gangemi, 1991; Wong and Watt, 1990)

2. Security in the OSI Model

There are five basic security services described by the ISO Security Architecture that can be implemented in any network. These services are authentication, access control, data confidentiality (i.e. secrecy), data integrity and nonrepudiation (ISO 7498-2-1988(E)). *Authentication* is the process of verifying the identity of a user or computer process acting on behalf of a user. *Access control* (i.e. authorization) is the process of enforcing the authorized use of or access to programs, processes and data by users or computer processes acting on behalf of users. *Integrity* implies that there has been no accidental or intentional unauthorized modification or destruction of data. *Confidentiality* implies that there has been no accidental or intentional unauthorized disclosure of data. *Nonrepudiation* is the process of protecting against denial of receipt or transmission of data by the receiver or sender respectively. (McNulty, 1994; Russell and Gangemi, 1991; Kirkpatrick, 1989) The goal of these five security services then is to prevent unauthorized disclosure, destruction or modification of data, or denial of system services in accordance with our definition of AIS security (DoN, 1991).

With the exception of the session layer, each notional layer in the OSI model can support one or more of these security services. Figure 2 illustrates the security services that each layer can support. The physical layer can support data confidentiality using end-to-end encryption protocols. The data link layer can also support data confidentiality by using link-by-link rather than end-to-end encryption protocols. The network and transport layers can each support authentication, access control, data confidentiality and data integrity. The presentation layer can support data confidentiality. The OSI application layer can support all five security services. (Buddenberg, 1995; Russell and

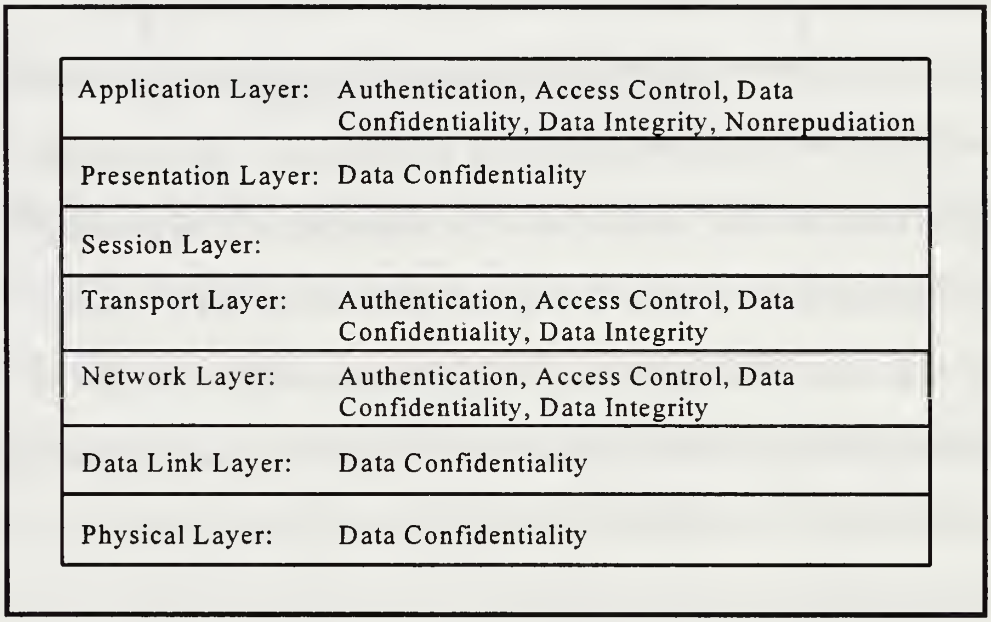

Application Layer:	Authentication, Access Control, Data Confidentiality, Data Integrity, Nonrepudiation
Presentation Layer:	Data Confidentiality
Session Layer:	
Transport Layer:	Authentication, Access Control, Data Confidentiality, Data Integrity
Network Layer:	Authentication, Access Control, Data Confidentiality, Data Integrity
Data Link Layer:	Data Confidentiality
Physical Layer:	Data Confidentiality

Figure 2. Security services in the OSI Reference Model. After (Kirkpatrick, 1989).

Gangemi, 1991; Kirkpatrick, 1989) It is important to note that the OSI security architecture only discusses the services each layer *can* support; it says nothing about which services each layer *should* support.

D. INTERNET PROTOCOLS

As mentioned in Section C, the OSI reference model is largely notional. Actual network technology and protocol implementations do not map neatly to the OSI conceptual model, yet actual implementations still provide essentially the same overall functionality. Subsection 1 below presents a brief discussion of the common TCP/IP protocol suite and how it maps onto the OSI reference model. Subsection 2 discusses computer security in the Internet model.

1. Network Functionality in the TCP/IP Protocol Suite

The TCP/IP protocol suite is composed of four conceptual layers: the network interface layer, the network or Internet layer, the transport layer and the application layer. Figure 3 illustrates how these layers map onto the OSI model. Although the figure shows a smooth mapping, one must realize that it is not necessarily a direct mapping; there is some functional overlap between the layers of the OSI model and those of the TCP/IP suite. A significant further difference is that protocols following the OSI reference model are constrained to communicate only with adjacent layers, whereas interlayer communications in the TCP/IP suite are allowed to bypass layers when necessary for efficiency, responsiveness and throughput (Stallings, 1994).

The network access layer in the TCP/IP suite combines most of the functionality of the OSI physical and data link layers. However, where the OSI model includes link-by-link error control at the data link layer, the TCP/IP protocol stack treats error control as an end-to-end problem in the transport layer. The network access layer is responsible for accepting IP datagrams from the network layer, encapsulating them in network

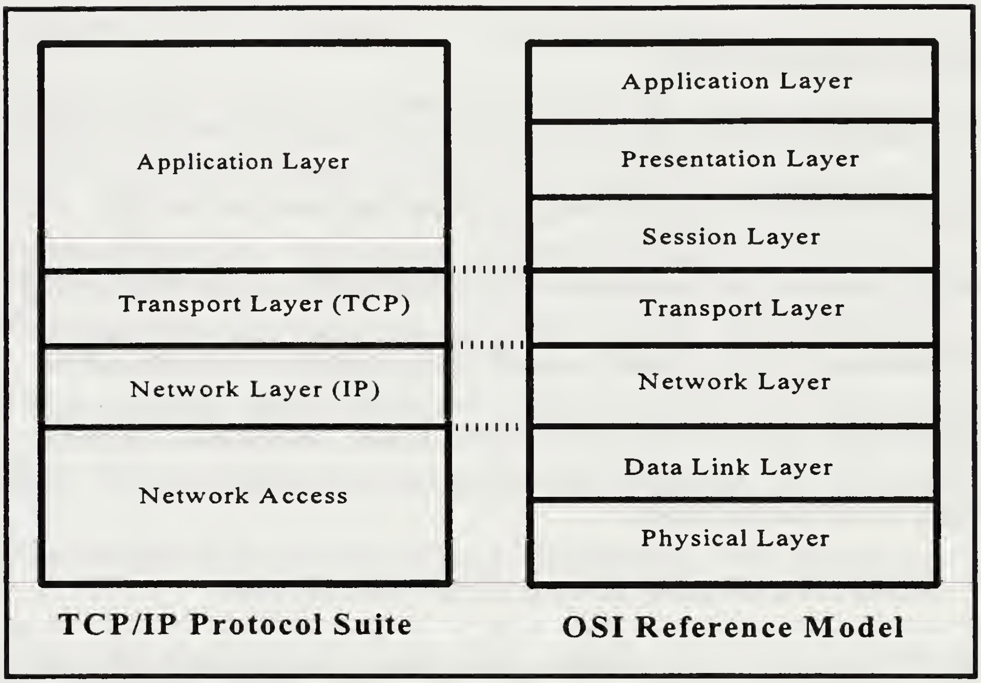

Figure 3. Comparison of the TCP/IP protocol suite and the OSI model. After (Varma, 1995).

specific frames and transmitting them over the network. The TCP/IP network access layer may consist of a simple device driver or it may implement its own complex data link protocol. (Comer, 1991)

The IP network layer (also known as the Internet layer) provides functionality similar to that described for the OSI network layer. However, the OSI model only describes functionality associated with a single network; it does not address internetworking as does the TCP/IP implementation. The IP network layer accepts packets from the transport layer, encapsulates them as Internet Protocol (IP) datagrams and passes the IP datagrams to the network access layer. (Comer, 1991)

As mentioned above, the TCP/IP transport layer provides end-to-end error control and other functionality similar to the OSI transport layer. Protocols in this layer include the connection-oriented Transport Control Protocol (TCP) (Postel, 1981) and the connectionless User Datagram Protocol (UDP) (Postel, 1980). The TCP/IP transport layer may also implement some OSI session layer connection management functionality in those instances where TCP is implemented. (Comer, 1991)

Finally, the TCP/IP application layer combines the functionality described for the OSI session, presentation and application layers. In the TCP/IP implementation, the application layer is responsible for determining whether the data is to be transmitted as a data stream or as data messages (Comer, 1991).

2. Security in the TCP/IP Protocol Suite

The TCP/IP protocol suite is a set of interoperable protocols for use in an open environment. For this and historical reasons, the current version (IPv4) protocols themselves are inherently insecure. The current TCP/IP protocols do not prevent someone at an intermediate node between source and destination nodes from capturing and viewing, modifying, destroying, spoofing or redirecting datagrams as they pass by on the network.

Many tradeoffs pertain when implementing security in practice. The primary focus for security in the TCP/IP model is at the application layer (Baker *et al.*, 1996). Implementing security in higher layers typically requires a smaller Trusted Computing

Base (TCB)[2] to enforce a particular security policy. Generally a smaller TCB facilitates scalability and reduces security related costs.

Despite the inherent insecurity of the current Internet protocols, they do support some security functions. For example, the transport layer incorporates an integrity check-sum in the transport control protocol (TCP) packets and data confidentiality can be supported at either the network or application layer using cryptography (Comer, 1991). Access control can also be supported at the application or network layer with address filtering or access control lists.

Security requirements have changed as the Internet has migrated from primarily a research and development tool to encompass a much broader user base. The IP Security Protocol (IPSEC) working group within the Internet Engineering Task Force (IETF) has developed cryptographic security protocol specifications to support combinations of access control, authentication, confidentiality and integrity at the network layer (Hughes, 1996; Atkinson, 1995).[3] These specifications have been incorporated into the next-generation Internet protocol (IPv6) (Bradner and Mankin, 1996; Huitema, 1996; Nierle, 1996). While the specifications are compatible with current IPv4 protocols, they are not widely used (Huitema, 1996).

[2] The Trusted Computing Base (TCB) is the "totality of protection mechanisms within a computer system — including hardware, firmware, and software — the combination of which is responsible for enforcing a security policy" (DoD, 1985).

[3] Refer to Section F later in this chapter for a discussion of the IETF. Information on the IPSEC working group and their work can be found at *http://www.ietf.org/html.charters/ipsec-charter.html*

E. ASYNCHRONOUS TRANSFER MODE

Asynchronous Transfer Mode (ATM) is an emerging transport and switching technology for use in broadband digital communications. It offers the possibility of transmitting a variety of media formats at speeds ranging from a few megabits per second to gigabits per second. A good overview of the technology is in (Partridge, 1994) and an in-depth discussion can be found in (Alles, 1991). This section serves only to introduce the reader to how ATM compares to the TCP/IP protocol stack and to provide an overview of security as it applies to ATM. A detailed study integrating ATM with LAN and WAN connectivity at NPS is documented in (Courtney, 1996).

1. Internet Protocol (IP) and ATM

IP and ATM are fundamentally different. IP is a connectionless technology whereas ATM is a connection-oriented technology. Current ATM research and implementations focus on combining existing IP protocols and immature ATM protocols. Figure 4 illustrates how ATM network services and protocols roughly map onto the TCP/IP protocol suite.

The physical layer in the ATM protocol model is still responsible for bit encoding, but the ATM layer combines the functionality of the OSI data link and network layers for converting data link and higher layer protocols (such as TCP and IP) into ATM specific protocols (Alles, 1991). The ATM Adaptation Layer (AAL) is responsible for segmentation and reassembly of the data stream. The AAL also handles multiplexing and cell loss detection responsibilities and retransmission of lost cells for connection-oriented service. (Varma, 1995)

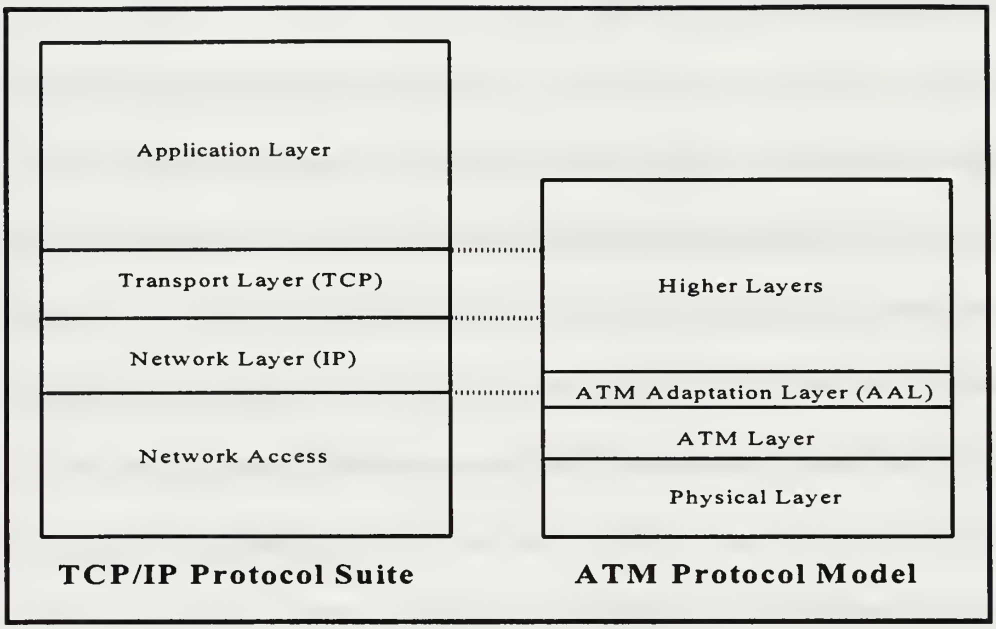

Figure 4. Comparison of the TCP/IP and ATM models. After (Varma, 1995).

2. ATM Security

a. Vulnerabilities

ATM can be implemented over a variety of physical media. ATM is envisioned as the next-generation broadband digital transmission method, usually run using Synchronous Optical Network (SONET) protocol over optical fiber. There may be a tendency to assume optical fiber is inherently secure. While it is true that some forms of attack are more difficult with fiber optic links (line tapping for example), attacks are not impossible.

There also may be a tendency to assume that "packet-sniffing" is not possible in an ATM network because ATM is a connection-oriented technology and cell switching occurs in hardware rather than software. This also is not true. ATM is vulnerable to many of the same types of attack that have befallen the Internet community, specifically eavesdropping (i.e. packet/cell-sniffing). Cell-sniffing was demonstrated by MCNC[4] using the same techniques used for IP packet-sniffing. MCNC connected a traffic analyzer to an ATM switch and captured login sequences (i.e. user IDs and passwords) traversing the network (Stevenson *et al.*, 1995). While this may not be a significant vulnerability for a local-area network (LAN) where all switching hardware is physically controlled by an organization, it can be quite significant for "logical" local-area networks (i.e. those sharing the same address space) that are spread over a wide area, and for wide-area networks (WANs) that use the public ATM network infrastructure which is not controlled by a single organization. (Pucher *et al.*, 1996; Stevenson *et al.*, 1995)

In addition to the cell-sniffing vulnerability, ATM is vulnerable to masquerading attacks. The ATM signaling protocol does not implement a way to ensure the authenticity of a user accessing a system. The calling party ID field in the SETUP signaling message is optional and is supplied by the calling party. The called party has no way of ensuring the accuracy of the information in the calling party ID field and therefore has no way of ensuring the authenticity of the calling party. (Stevenson *et al.*, 1995)

ATM is also vulnerable to denial-of-service attacks. The RESTART signaling message can be used to deallocate bandwidth resources associated with any

[4] MCNC is a "nonprofit corporation that develops and applies advanced electronic and information technologies for business, university research, government, and education for economic development" within North Carolina. MCNC is the corporation's legal name, not an acronym. The MCNC Mission Statement is available at *http://www.mcnc.org/HTML/welcome.html*

Virtual Circuit (VC) specified in the message, including every active VC for another user. The standard signaling protocol does not require proof of identity for the user issuing the request. (Stevenson *et al.*, 1995)

b. Proposed Security Services

Recognizing the vulnerabilities in ATM, researchers have recently begun investigating ways to include security services in the lower-level ATM protocols. Researchers now recommend that authentication, data integrity, data confidentiality and cryptographic key distribution be implemented in lower-level ATM protocol standards. Proposed ATM security standards recommend that data integrity be supported at the AAL level and that authentication, cryptographic key distribution and confidentiality be supported at the ATM layer. It also is recognized that these security services are required not only for the data stream, but for signaling and network management messages as well. (Chuang, 1995; Peyravian and Herreweghen, 1995)

Recommending inclusion of data integrity at the AAL layer rather than the ATM layer is due to overhead limitations in the individual data cells. An ATM cell is only 53 bytes long: 48 bytes for the data payload and 5 bytes for header information (Peyravian and Herreweghen, 1995; Varma, 1995; Partridge, 1994). Increasing the header size to include an integrity checksum would necessarily reduce the cell's data payload. Such a change might result in an unacceptable reduction in data throughput (Peyravian and Herreweghen, 1995).

Recommendations for including authentication, data confidentiality and cryptographic key distribution at the ATM level are described for three scenarios in

(Peyravian and Herreweghen, 1995). The first scenario calls for security services from one ATM end-node to another ATM end-node across a public (or private) network. The second scenario calls for security services from an ATM end-node to an ATM switch on the border of a public (or private) network. The third scenario calls for security services between border switches residing on separate public (or private) networks. These scenarios are shown in Figure 5. (Peyravian and Herreweghen, 1995)

The specification defining the proposed security standards for these scenarios is scheduled to be published by The ATM Forum in mid 1997 (Sheth, 1996). (A discussion of The ATM Forum and their efforts is included in the next section.) However, previous experience with the closed nature of the ATM Forum indicates this estimate may be optimistic. As a result, such proposed standards and corresponding delays in implementation will not be of immediate use to the NPS ATM LAN. Furthermore, even if the ATM Forum remains on schedule, the specifications are likely to be incomplete as closed-door security reviews typically overlook crucial issues that are only revealed through widespread critique and discussion. That ATM security problems are not widespread at present is probably due to the immaturity of ATM technology, incompatibility among ATM switches even during normal operations and the lack of widespread knowledge or use of ATM. These problems are documented in detail in (Courtney, 1996). Thus all aspects of ATM security need to be considered suspect for the foreseeable future. Interested readers should consult the ATM Forum to remain up-to-date on current developments.

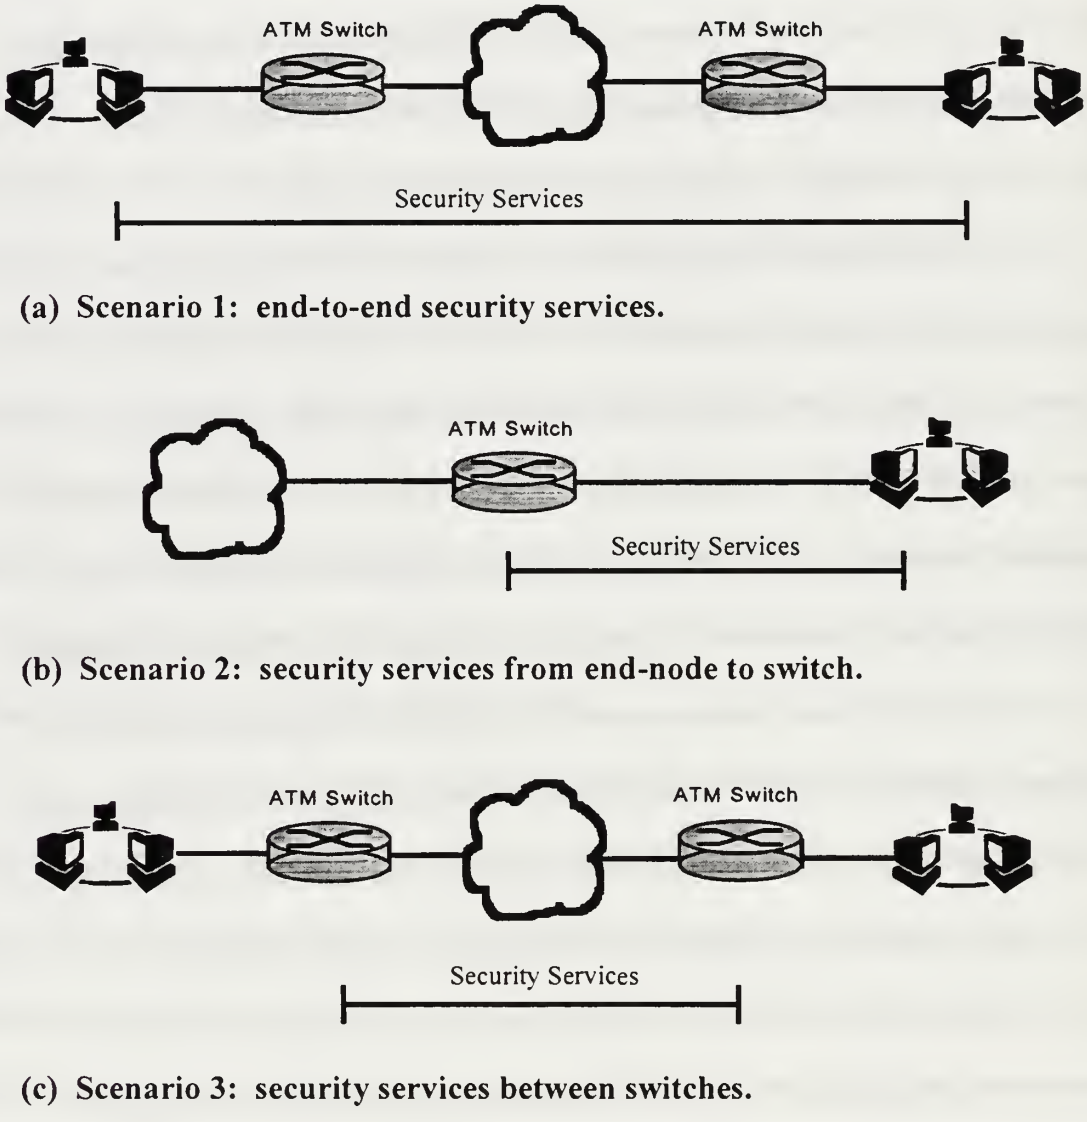

(a) Scenario 1: end-to-end security services.

(b) Scenario 2: security services from end-node to switch.

(c) Scenario 3: security services between switches.

Figure 5. Categories of forthcoming proposed security service implementations in ATM networks. After (Peyravian and Herreweghen, 1995).

F. RELATED WORK

This section identifies related research being conducted in internetworking and computer security. Subsection 1 examines the impact of an important network computing security thesis at NPS. Subsection 2 highlights the research being done by the Information Infrastructure Research Group (IIRG) at NPS. Subsection 3 provides an introduction to the Internet Engineering Task Force (IETF) and its work. Subsection 4 introduces the ATM Forum and its research. Subsection 5 briefly explains the Information Wide-Area Year (I-WAY) high-performance-computing project. A short discussion of (IPv6) is presented in Subsection 6. This section concludes in Subsection 7 with a description of some additional related NPS research projects.

1. Unix System Security Analysis

Unix Security: A Penetration Analysis of Navy Computer Systems (Rich, 1992) examines and documents vulnerabilities associated with a number of U.S. Navy networked Unix computer systems connected to the Internet. This thesis examines six well-know, well-documented Unix security vulnerabilities and demonstrates the ease with which those vulnerabilities can be exploited. This thesis documents the results from security sweeps of four Navy commands. While this thesis does indicate permission was obtained from each command prior to the sweep, it does not indicate if the responsible system administration managers were notified prior to the sweep. Such secretive "attacks" have been permitted at NPS. If the purpose of such secretive "attacks" is to discover existing system vulnerabilities system administrators need to be notified. Even if the purpose of such secretive "attacks" is to test the effectiveness of system administrator training in a realistic setting, this practice unnecessarily burdens already heavily tasked system administrators since they can end up chasing the proverbial "wild goose" in their attempts to isolate and track down the intruder.

The strength of this thesis is in its recommendations for a simplistic Unix security model, a model that was largely ignored until after the December 1995 NPS security incident. Unfortunately access to the thesis is limited to the DoD and DoD contractors simply because it documents computer system security vulnerabilities even though the vulnerabilities are well-documented in other sources, (Curry, 1990) for example. Limiting access to reports such as this is a misguided attempt at improving system security by keeping quite and hoping intruders are ignorant about common vulnerabilities. Limiting access to this particular thesis no doubt contributed to the successful NPS break-in in December 1995.

2. Information Infrastructure Research Group (IIRG)

This thesis is one of several master's theses being written as a result of cooperative research into ATM and other broadband networking technology applications by the Information Infrastructure Research Group (IIRG) at NPS (IIRG, 1996). The following sections provide a brief overview of other work being done, identifying how networked computer security applies to each area of research.

a. Automated Local and Global Network Monitoring

Internetworking: Automated Local and Global Network Monitoring (Edwards, 1996) evaluates existing public domain network monitoring tools. The tools are evaluated on IP and ATM LANs and WANs.

Network monitoring includes the analysis of message transmission time and routing. These include some of the same types of tools used by Cliff Stoll and his colleagues at the Lawrence Berkeley Labs to track down a band of West German hackers in the early 1980's (Stoll, 1989). Just as Stoll discovered during his hacker hunt, this thesis recognizes the complexity of large-scale internetworks and the problems associated with administration and monitoring of those networks.

A variety of both commercial and public-domain tools exist to aid system administrators. Unfortunately, commercial tools are prohibitively expensive. On the other hand, public-domain software tools are often difficult to understand and use. The goal of this thesis is to integrate various public domain software tools into a single automated easy-to-use network monitoring package that can provide network administrators the network status and performance information they need.

Security considerations summarized in this thesis include using automatic scripts accessible via home pages to remotely monitor network status. The monitoring scripts can also automatically e-mail or page system administrators to inform them of network problems.

b. ATM Local-Area Network (LAN)

Internetworking: The Naval Postgraduate School (NPS) ATM Local-Area Network (LAN) (Courtney, 1996) discusses ATM LAN configuration and compatibility issues and documents the steps taken to develop a simple ATM LAN at NPS.

The objective of this project is to create, test and build an electronic information infrastructure at NPS based on ATM cell relay that will enable research work in tele-education, telescience experiments and digital interactive multimedia. This thesis lays the groundwork for future ATM work by sending real-time audio, video and graphics across campus and around Monterey Bay over ATM circuits.

The primary security considerations highlighted in this thesis are system availability issues associated with multi-vendor hardware and software incompatibility. This is a prototype project and therefore known security problems have not prevented careful implementation, testing and evaluation. It is recommended that the security vulnerabilities and issues associated with ATM (discussed earlier in this chapter) will be addressed prior to any operational implementation.

c. Multicast and ATM Network Prerequisites for Distance Learning

Internetworking: Multicast and ATM Network Prerequisites for Distance Learning (Tamer, 1996) investigates simple and cost-effective ways to configure and use the Internet multicast backbone (MBone) and related software tools to provide adequate audio and video quality over the Internet as well as ATM networks where available. This research develops an inexpensive and easy-to-learn method for implementing and operating an MBone classroom facility. This thesis outlines the steps required to configure MBone audio/video recording tools to allow educational institutions with limited budgets (most notably K-12 schools) the opportunity to take advantage of the potential educational benefits offered by the MBone.

This thesis also investigates the performance of multicast transmissions over ATM. Current MBone tools provide adequate (although low frame rate) video transmissions over the Internet, but performance is limited by the available bandwidth.

Although ATM is a young technology, it's high bandwidth and low latency offers the promise of near-real-time broadcast quality audio/video transmissions.

Security is not addressed in this thesis. However, educators who are considering an MBone implementation must examine such issues as hardware access, system availability, acceptable use of the resources and training to ensure that resources are properly used, so as not to degrade the performance of the regional and global MBone.

d. Live Multicast Audio/Video Distribution

Internetworking: Using Global ATM Networks for Live Multicast Audio/Video Distribution (Erdogan, 1996) documents implementation of the MBone over a regional Frame Relay wide-area network (WAN).

The Monterey BayNet regional Frame Relay WAN links various K-12 schools, libraries, research and other educational partners around the Monterey Bay area, but does not effectively take advantage of multicasting technology for distance learning projects. This thesis documents an inexpensive MBone implementation for distance learning over Frame Relay. Despite the recommended requirement for multicast routers for multicasting in a frame relay network, this work also illustrates the feasibility of implementing multicast service over Frame Relay without using multicast routers.

This thesis also does not discuss security issues *per se*, but the same issues discussed above for (Tamer, 1996) apply to this project. System availability will likely be the primary concern for any educational implementation of the MBone. Proper configuration and use of the MBone tools will ensure continued availability. An emphasis on training is required. Additionally, this work parallels the monitoring work documented in (Edwards, 1996) in that automatic scripts accessible via home page are used to monitor MBone status. Although the software tools used in the scripts typically operate with root permissions, software analysis shows that public access to these scripts does not represent a confidentiality vulnerability. See Chapter IV for a general discussion of these issues.

e. Economical Storage and Retrieval of Digital Audio and Video for Distance Learning

Internetworking: Economic Digital Storage and Retrieval of Digital Audio and Video for Distance Learning (Tiddy, 1996) compares and evaluates existing digital storage methods for audio and video.

While the Internet already provides for transmission of near-real-time audio and video over the Internet using either unicast or multicast messages, current audio/video tools for viewing archived files require the entire file be downloaded to the local workstation before it can be viewed. For most users this is infeasible with large files. Therefore this thesis compares and documents the efficiency of existing transfer modes, file formats and data compression methods for on-demand audio and video over the Internet.

This thesis also does not specifically address security issues. However, similar to (Edwards, 1996) and (Erdogan, 1996), this project implements scripts that are accessible via home pages. The scripts receive requests to transfer archived audio/video files to remote users. Anytime a server receives input from external users (especially path information for file manipulation) there is the possibility that it can be subverted (Stein, 1996).

f. Implementing IPv6 in the Marine Corps Tactical Data Network

Internetworking: Technical Strategy for Implementing the Next Generation Internet Protocol (IPv6) in the Marine Corps Tactical Data Network (Nierle, 1996) describes an Internet Protocol (IP) addressing plan for use in the Marine Corps' Tactical Data Network (TDN). This thesis discusses the current Marine Corps internetworking requirements but also recognizes the weaknesses associated with the existing Internet Protocol version 4 (IPv4) and the yet-to-be--demonstrated security features of the next-generation Internet Protocol version 6 (IPv6).[5] Nierle presents a tactical addressing plan for the TDN that works with IPv4 and facilitates smooth migration to IPv6 as the technology matures.

Security considerations summarized in this thesis include the need for data security as well as physical security of networking hardware. This thesis also describes the enhancements of IPv6 over IPv4 concerning native security support at the Internet (IP) layer and improved mechanisms for controlling the quality of service (QoS).

3. Internet Engineering Task Force (IETF)

The Internet Engineering Task Force (IETF) is an open international community composed of network designers, researchers, vendors and users. The IETF is the engineering and protocol development arm of the Internet. The IETF is organized into a

[5] IP version 5 was assigned to another protocol which was subsequently unused (Huitema, 1996).

number of focus areas, each with its own specific charter. Each focus area is composed of a number of working groups. The current IETF focus areas and the Security Area working groups are shown in Table 5. (IETF, 1996)

Chapters V and VII in this thesis describe the Kerberos authentication and authorization system and its implementation at NPS. Kerberos is a secure network access mechanism that relies on a trusted host to provide authentication services without sending passwords over the network in the clear. Work relating to Kerberos and the use of public-key cryptography for the initial Kerberos authentication is covered by the Common Authentication Technology (CAT) working group. This group focuses on providing protocol-independent distributed security services. Their work focuses on security

- Applications Area
- Internet Area
- Network Management Area
- Operational Requirements Area
- Routing Area
- Security Area
    - Authenticated Firewall Traversal (aft)
    - Common Authentication Technology (cat)
    - Domain Name System Security (dnssec)
    - IP Security Protocol (ipsec)
    - One Time Password Authentication (otp)
    - Public-Key Infrastructure (X.509) (pkix)
    - Transport Layer Security (tls)
    - Web Transaction Security (wts)
- Transport Area
- User Services Area

Table 5. IETF focus areas and Security Area working groups (IETF, 1996).

implementation tasks and interfaces rather than on integrating security data elements into specific protocols. Mailing list information for the CAT working group follows.

IETF CAT Working Group Mailing List Information:

General discussion:	*cat-ietf@mit.edu*
To subscribe:	*cat-ietf-request@mit.edu*
Archive:	*ftp://bitsy.mit.edu/cat-ietf/archive/*

4. The ATM Forum

The ATM Forum is an international non-profit organization composed of over 700 member companies. The organization's objective is to accelerate the use of ATM technology through industry cooperation in protocol and standards development, as well as education of the networking and telecommunication community on the capabilities of ATM (ATM, 1996). The security working group within the ATM Forum was established in July 1995 for the purpose of defining ATM security specifications (Rendleman, 1995). The group is not expected to publish a final version of *ATM Security Specification (SS) 1.0* until mid-1997 (Sheth, 1996). Work completed by the ATM Forum is available for a fee to non-members. Work in progress typically is not available for review or comment. Occasionally ATM Forum development efforts are performed openly in cooperation with IETF working groups.

5. Information Wide-Area Year (I-WAY)

The Information Wide-Area Year (I-WAY) is an experimental high-performance internetwork linking approximately 30 of the country's fastest computers and advanced visualization environments on approximately 11 different regional ATM networks.[6] The

[6] Information about the I-WAY is available at *http://www.iway.org*

I-WAY supports both TCP/IP over ATM and native ATM-oriented protocols. The supercomputers linked by the I-WAY provide the resources required to run simulations in a variety of virtual environments. To control remote access to these resources, the I-WAY uses the Kerberos authentication and authorization protocol. The Kerberos implementation at NPS began as a class project to establish connectivity with other I-WAY participants during the Supercomputing '95 conference held in San Diego California in December 1995. (Karin, 1995)

6. Internet Protocol version 6 (IPv6)

The existing Internet Protocol has a number of limitations both in its addressing structure and with security implementation that requires improvements. Internet Protocol version 6 (IPv6) is a designed to replace the current Internet Protocol version (IPv4). (Deering and Hinden, 1995) The new version remains a connectionless protocol but does define a number of improvements over IPv4. IPv6 expands the IP address size from 32 bits to 128 bits and also improves multicast message routing. IPv6 implements flow labeling, authentication and privacy capabilities. The flow labeling capability allows for specific quality of service (QoS) or "real-time" requests. The IPv6 specification also defines extensions to support authentication, data integrity and (optional) data confidentiality. One active area of work is getting vendors and users to implement these extensions. (Huitema, 1996; Nierle, 1996; Deering and Hinden, 1995)

7. Additional Related NPS Research

a. Integrating the Defense Message System and the Global Broadcast System

Tactical DMS: A Global Broadcast Service Option (Morales, 1996) describes a means of integrating the Defense Message System (DMS) and the Global Broadcast System (GBS). It also pertains to multicasting work being done by the IIRG. This thesis discusses the security and multicast advantages present in IPv6 and examines many of the same tactical issues presented in (Nierle, 1996).

b. Trust and Interoperability in a Networked Environment

Ensuring a C2 Level of Trust and Interoperability in a Networked Windows NT Environment (Lucas, 1996) provides an in-depth discussion of the advantages, disadvantages and applicability of the Trusted Computer System Evaluation Criteria (TCSEC) (DoD, 1985) for today's military networks. This project describes the broad scope of computer security in today's computing environment and emphasizes the need for "trusted systems" to ensure confidentiality, integrity and system availability. This thesis presents a case study of a military network using Windows NT version 3.51 and evaluates how effectively this network operating system meets the requirements of the TCSEC.

G. SUMMARY

This chapter provides the background information necessary for the reader to understand what computer security is and why computer security is important. Department of the Navy AIS security considerations are discussed and the (often hazy) distinction between the terms threat, vulnerability, risk, and control are clarified.

Network communication and security service functionality concepts within the ISO OSI reference model are discussed. Network communication and security functionality associated with the TCP/IP protocol suite and ATM protocols are also briefly discussed. Some of the vulnerabilities associated with ATM are highlighted. The chapter concludes with a broad discussion of related internetworking and security research.

III. PROBLEM STATEMENT

A. INTRODUCTION

There are two primary variables that affect the security of any information system: technology and people. Technology issues must address what security is, how best to implement existing capabilities, and how to integrate emerging technologies while remaining secure. People issues must address how much security is required, as well as who, when, why and from where access to the computer system/network will be allowed. The two perspectives frequently overlap when new technology is applied to support security. Section B defines the two primary problems this thesis addresses: (1) understanding the scope and need for computer system/network security for unclassified systems and (2) countering the threat of unauthorized system/network access resulting from open and interoperable networks. Section C briefly describes the solutions presented in this thesis.

B. PROBLEMS

1. Understanding Networked Computer System Security

Networked computer system security is a complex problem, one that is no longer solely restricted to military or national security and associated *classified* computer systems and networks. In 1996 there were over 250,000 attempted break-ins to *unclassified* systems within the Department of Defense (DoD); two-thirds of those were successful (Zuckermann, 1996). Computer/network security now encompasses a much wider scope than just confidentiality and basic access controls. Accelerating trends in networking requires a shift in how managers and users view security. Computer and network professionals must understand the scope of security and recognize the need for security for even their unclassified systems.

One tenet of the IIRG is that technical problems have technical solutions, and people problems have people solutions. Difficulties often arise when problems and solutions are mismatched. This thesis attempts to be very clear in distinguishing whether problems and solutions are technology-related or people-related.

AIS security is primarily a management issue in that technical, administrative and procedural controls implemented to secure information systems cost money. Regardless of the size of an organization, managers control the money. Managers must set priorities, make decisions and allocate scarce organizational resources. Therefore managers must ultimately determine the level of security required for any given system, and managers must accept responsibility for the consequences if information security proves inadequate. Of course, managers are foolish if they make decisions without first examining the issues involved in making those decisions.

The main challenge with securing a networked automated information system is recognizing that threats to AIS security exist and admitting that the system in question is not immune to those threats. The specific problem is in articulating the need for AIS security in the form of a quantitative security policy that is acceptable to all stakeholders. Managers must then take the necessary implementation steps, allocate the necessary resources and provide the necessary support to carry out that policy. Proper planning, analysis and implementation can ensure a sufficiently secure system.

2. Open Systems

The emphasis on open and interoperable networks has left many systems vulnerable to a wide variety of attacks. In a stand-alone computing environment, threats to AIS security are relatively easy to control. However, once computers are connected together to form a single integrated network, which is then connected to other integrated networks form internetworks, the risks associated with previously existing threats are greatly increased and new threats are introduced. Additionally, integrating any new technology into an existing infrastructure can be problematic and introduce previously unforeseen security threats. (Fraser, 1996; Fitzgerald, 1993; Rich, 1992)

Transmission of static passwords in plain text over the Internet has been one of the most widely publicized vulnerabilities to network security (CERT, 1994; NIST, 1994; McNulty, 1994). Mr. F. Lynn McNulty, the Associate Director for Computer Security at the National Institute of Standards and Technology (NIST), has testified before Congress that "traditional user authentication by means of re-useable passwords does not provide strong security in today's networked environment -- with or without encryption" (McNulty, 1994, pg. 59). The use of one-time passwords (such as the S-Key protocol (Haller, 1995)) or other secure remote access mechanism (such as the Kerberos protocol (Miller, *et al.*, 1987)) has been recommended to protect against unauthorized access to computer systems connected to the Internet (NIST, 1994; McNulty, 1994; Bennington, 1991).

As a graduate education and research institution, the Naval Postgraduate School (NPS) relies heavily on computer networking and internetworking. NPS must maintain open connectivity with local and remote research partners if it is to accomplish its mission. NPS must also take the necessary steps to protect its systems and the information stored on those systems from unauthorized access. This is a precarious balance. Although an open environment promotes collaborative research, it is also an inviting target. Therefore security is necessary if for no other reason than to maintain connectivity and to protect research information.

C. EXPERIMENTAL SOLUTIONS

This thesis examines many of the issues that must be addressed when assessing the need for computer system/network security. The issues discussed are used as a basis for a site security survey for the unclassified integrated IP/ATM LAN/WAN in the Systems Technology Lab (STL) at NPS. This assessment is merely a starting point; recommendations for additional steps needed to improve the state of security are included.

To improve remote plaintext password protection we implement and examine the feasibility of the Kerberos authentication protocol. As with the security survey, this is merely a starting point. There remains much work to be done. Nevertheless the results presented in this thesis are promising and productive.

D. SUMMARY

There are usually two aspects to any problem: a social aspect and a technological aspect. This chapter explains the problem of computer system/network security from these two perspectives. Computer system/network security is first a social issue in that it affects people. Managers and users alike often do not understand networked computer system security or the need for it, yet managers are responsible for ensuring security is adequate yet not restrictive. Computer system/network security is also a technological issue for two reasons: technology often exacerbates the social issues, and management policy concerning system/network security is often enforced through the use of technology. Section B briefly describes the two primary problems associated with system/ network security from these two perspectives and Section C lists the solutions presented in this thesis to address these problems.

IV. INTEGRATED IP/ATM LAN/WAN SECURITY ISSUES

A. INTRODUCTION

Automated information system (AIS) security is primarily a management challenge (Walker, 1991). This chapter discusses some of the issues managers face when deciding how to enhance the security of their computer system/network. Security controls are not discussed in detail here since they are examined extensively in a number of other sources (Cobb, 1992; DoN, 1991; Wong and Watt, 1990; Palmer and Potter, 1989; Wood *et al.*, 1989). While the decision outcomes concerning security vary from system to system, the issues are largely technology-independent and must be addressed for each system. Section B discusses the administration issues of AIS security policies, plans and procedures. Environmental issues relating to natural and human events are discussed in Section C. Software and data security issues are discussed in Section D. Security issues associated with telecommunications and system access are discussed in Section E. Finally, issues concerning security implementation are discussed in Section F.

The issues highlighted in this chapter are used as a basis to assess the status of security associated with the integrated IP/ATM LAN/WAN at NPS. The results of that assessment are presented in Chapter VI.

B. ADMINISTRATION ISSUES

How much security is required? What types of controls are required? These questions are not trivial and must be answered before scarce resources are allocated to implement any security controls. Before these questions can be answered, system administration managers must identify the purpose and acceptable uses of the information system in question, and they must identify the individuals who will be responsible for implementing and administering system security. Furthermore, managers need to understand the environment in which they are operating and the threats such an environment presents. Management decisions must then be documented and disseminated to as wide an audience as possible. Users, system administrators and other managers alike must be made aware of these decisions. This section discusses the need for managers to address such issues prior to implementing security controls. Subsection 1 discusses system availability requirements, Subsection 2 discusses system security and acceptable use policies (AUPs) and Subsection 3 discusses contingency planning.

1. System Availability Requirements

Protecting against denial of system service is one of the basic goals of computer system/network security. (See the definition of AIS security in Chapter II.) It also is probably the most difficult goal to achieve. Like many other aspects of system/network security, protecting system services is itself a multifaceted problem. Nevertheless system availability must be addressed because an information system is useless if users cannot get or use the system resources they need.

One of the primary issues that will determine the amount of security required for an information system is the importance of the system and the information it processes. System availability requirements (i.e. maximum allowable down time) will vary according to this assessment and must be clearly understood. A critical system may have a 99% (or greater) availability requirement (e.g. down time not to exceed 15 minutes in a 24 hour period), whereas a secondary or noncritical system may be able to accommodate much longer down times. (Buddenberg, 1995; Russell and Gangemi, 1991; Wong and Watt, 1990)

Of related interest is the ongoing lack of any system availability requirement for NPS, recently demonstrated by a 3-week "security shutdown" of all networked computing systems in late 1995. This occurred despite the fact that one of the major goals of the NPS Automated Data Processing Security Program (NPS, 1992) is to "minimize the denial of ADP services to authorized users which might occur as a result of unauthorized access, operator error, power disruption, malicious acts, etc." (NPS, 1992, pg. 2). Repercussions of such shutdowns jeopardize the operations, mission and long-term viability of the school.

Threats to system availability are not limited strictly to the environmental, telecommunications or system-access threats discussed later in this chapter. Restricted computing center operating hours, inadequate help facilities (human or automated), and even inadequate memory, disk space or other resources can also lead to denial of service. While a complete discussion of how to determine system availability requirements is beyond the scope of this work, managers must carefully consider such requirements before any security controls are implemented. Without a careful analysis of system availability requirements, managers are unable to evaluate the cost effectiveness of the various controls. (Buddenberg, 1995; Russell and Gangemi, 1991; Wong and Watt, 1990)

2. System Security and Acceptable Use Policies (AUPs)

Before managers can determine how much and what type of security controls to implement for any given AIS, they must determine their goals and priorities concerning security for the organization and that system. This determination constitutes an organization's AIS security policy which clearly outlines "the set of rules, principles and procedures that regulate how an organization manages, protects, and controls . . . computer resources and the information they contain" (Cobb, 1992, pg. 111). The system security policy needs to clearly describe organizational priorities and philosophy concerning AIS security. This policy must be explained through training and distributed freely throughout the organization. (Fraser, 1996; Cobb, 1992; Russell and Gangemi, 1991; Wong and Watt, 1990)

A security philosophy will fall somewhere on the continuum between "denying all" system/network access and services and "allowing all" system/network access and services (Fraser, 1996; Russell and Gangemi, 1991). This is the fundamental dichotomy in the protection against denial of system services. A "deny all" policy will almost certainly eliminate unauthorized access but will also undoubtedly deny use of required services to otherwise legitimate users. Conversely an "allow all" policy will give users access to all required services but will also make those same services vulnerable to a denial-of-service attack. Therefore, policy decisions will reflect trade-offs: trade-offs between the available computing services offered and the amount of security provided; trade-offs between the ease of system use and the sophistication of security controls; and trade-offs between the AIS security costs (both tangible and intangible) and the risk of loss if a security incident occurs (Fraser, 1996).

The organization's acceptable use policy (AUP) for information resources is closely related to system availability and may be an integral part of the system security policy (Fraser, 1996). Automated information systems represent a large investment for most organizations and are usually employed in a multi-user environment. In such an environment, there is always the temptation for individual users to abuse or otherwise use the system resources either unproductively or inefficiently. If resources are being used for nonproductive or illegitimate purposes, they are unavailable for other users with legitimate and productive computing to accomplish. Borrowing a concept from economic theory, if users are not held personally responsible for their use of the resources, it is conceivable that the system might suffer a "tragedy of the commons." Such a situation occurs when communal resources are depleted through overuse or abuse by individuals who do not have a personal stake in the resource (Heyne, 1994). From a security standpoint, unauthorized use of information systems may present additional or unanticipated security threats to the system. Therefore managers must establish explicit acceptable use policies (AUPs). These administrative policies seek to ensure that the computing resources are used effectively and efficiently. Although NPS does have a campus-wide AUP (NPS, 1995), it shares the same problems of effectiveness and enforceability associated with any administrative policy. Users are not explicitly instructed or briefed on the AUPs nor are they required to sign a statement affirming acceptance of the policies.

3. Contingency Planning

No matter how much money an organization spends on IS security or how many controls are in place, disasters and security incidents can still occur. The actions an organization will take following such an occurrence must be planned and documented well in advance; it does little good (indeed it harms legitimate users) to lock one's front door after the thief has come and gone. These contingency plans must contain specific actions an organization will take, not only to react to a security incident or disaster but to recover from it as well. (Fraser, 1996; Fitzgerald, 1993; Cobb, 1992) The recovery actions outlined in the contingency plans will depend largely on the system availability requirements discussed above, therefore it is imperative that managers identify those requirements first (Palmer and Potter, 1989). NPS directives mandate that a contingency plan "be developed for each AIS where a disruption of service would have a critical impact on mission accomplishment" (NPS, 1992). Such a plan dos not exist and is not in wide use at NPS (Franklin, 1996; Norman, 1996).

Contingency planning is closely coupled with risk analysis and should include procedures for recovering from *all* threats to system security, not simply those presented by the environment. While it is important to have procedures in place in the event of a fire, flood, or power outage, it is equally important to plan for unanticipated hardware or software failures (such as a hard disk crash or software bug) and for malicious acts (such as a network virus or unauthorized network access). The more familiar personnel are with the contingency procedures, the faster the response is likely to be, quite possibly reducing the impact of the incident. (Fraser, 1996; Cobb, 1992; Wong and Watt, 1990; Palmer and Potter, 1989)

4. Roles and Responsibilities

AIS security roles and responsibilities for all personnel — managers, administrators, technicians and users — must be clearly documented and the affected personnel must be aware of (and intimately familiar with) the detailed requirements for their respective roles and responsibilities. AIS security responsibilities are often assigned to system administrators in addition to other roles and responsibilities. While this is not in itself unsatisfactory, managers must ensure that appropriate priority is given to the security responsibilities. As seen in Table 6, AIS security is a complex field with a great many roles and responsibilities that must be accomplished. Although some of these responsibilities will not apply to all organizations, it is easy to understand why, even in a small organization, AIS security requirements may quickly overwhelm even the most efficient system administrator. This is especially since most system administrators are already heavily tasked. If AIS security truly is a priority, managers must ensure adequate resources are allocated to satisfy that priority. (Fitzgerald, 1993; Cobb, 1992; Wong and Watt, 1991) NPS AIS "corporate" security roles and responsibilities are outlined in (NPS, 1992). Additional roles and responsibilities may be assigned on an individual departmental basis.

C. ENVIRONMENTAL ISSUES

Effective and efficient security controls cannot be implemented without a clear understanding of the physical environment in which an AIS is operating. Automated

- corporate IT security policy
- security of data center, communication network and terminals
- liaison with computer users
- production and enforcement of security standards and procedures
- Data Protection Act compliance
- risk analysis and monitoring
- implementation and administration of access control equipment and software
- control of encryption and authentication devices
- contingency planning and computer insurance

Table 6. IS security roles. From (Wong and Watt, 1991, pp. 128-129).

information systems are susceptible to a great many environmental hazards broadly classified as either a product of nature (e.g. earthquakes or floods) or a product of humans (e.g. intentional or accidental damage or destruction of equipment). Some environmental threats (unstable electrical power and fire, for example) can be caused either by nature or by man, but in such cases it usually does not matter how such disasters are caused since the end results are identical. Therefore the controls to counter these threats are likely to be identical. Natural threats are briefly discussed in Subsection 1 and human-induced threats are discussed in Subsection 2.

1. Natural Disasters

While often improbable, natural disasters are still quite possible. Threats in this category include not only familiar threats such as fires, floods, unstable electrical power and static electricity shocks, but also such uncommon threats as earthquakes, tsunamis and acts of war. Of course each site will have its own unique collection of threats. Controls in this category include but are not limited to: sprinkler systems, surge protectors or uninterruptible power supplies, anti-static carpeting, placing hardware on raised platforms, plastic sheets or covers to protect hardware from falling water, and temperature/humidity monitors to warn against excessive heat or humidity. (Palmer and Potter, 1989; Wood *et al.*, 1987) Minimum measures for protection against environmental threats are outlined in Appendix B of (NPS, 1992).

2. Human Disasters

Human disasters can be either accidental or intentional. Controls for intentional human threats typically focus on physical access security to protect against outside intruders. Hardware is locked up or access is otherwise controlled using a number of different mechanisms, ranging from simple escort procedures, access control lists and identification badges to more sophisticated equipment such as magnetic strip smart cards or biometric identification devices (e.g. finger print readers or retinal scanners). However, most human induced disasters are caused by internal personnel; authorized system users are far more likely to cause damage to the system than any external threat (Baskerville, 1989).

While disgruntled employees are a legitimate threat to system security, most damage caused by insiders is unintentional. The disgruntled employee threat can be minimized by good personnel security and proactive leadership. Accidental damage can be minimized through adequate administrative policies, consistent enforcement of those policies and adequate training. Additionally, software can provide simple automatic monitoring capabilities with system status results available via home page and electronic mail (Edwards, 1996; Erdogan 1996).

D. SOFTWARE AND DATA ISSUES

As costly as computer hardware may be, the most valuable asset associated with a computer system/network is usually the data and information it processes. This section discusses software and data security issues. Configuration management issues are discussed in Subsection 1; integrity and confidentiality issues are discussed in Subsection 2; the need for system backups is in Subsection 3.

1. Configuration Management

Configuration management is the process of "identifying, controlling, accounting for, and auditing all changes made to the baseline TCB [trusted computing base], including hardware, firmware, and software . . ." (Russell and Gangemi, 1991, pg. 145). As implied in the definition above, the term "configuration management" is often reserved for use with trusted systems (e.g. systems that process classified information), yet it is equally applicable to those systems processing and storing only unclassified or nonsensitive data. Improper system configuration or imperfect integration of application software can weaken system security. New application software releases can introduce additional security holes (CERT, 1996a; CERT, 1996b). Removing old or unneeded software not only simplifies file management but also avoids possible configuration conflicts with new software. Outdated operating system software with known security holes also must be updated with the latest software patch and old versions deleted from the system. System administrators must remain informed and take the necessary steps to minimize newly discovered threats and vulnerabilities of system and application software.

(Fraser, 1996; Curry, 1990; Rich, 1992) General configuration management requirements for NPS are outlined in (NPS, 1992).

2. Integrity and Confidentiality

Integrity and confidentiality of stored data can be protected using cryptographic checksums and bulk file encryption, but the primary means for protecting data integrity and confidentiality while it is stored and processed on an AIS is through data access controls. (Data integrity and confidentiality during transmission over network connections is discussed in Section E.) While physical access controls *indirectly* protect against data modification, destruction and disclosure by disallowing unauthorized access to the hardware components, data access controls *directly* protect against data modification, destruction and disclosure by disallowing unauthorized access to the software and data files.

Sensitivity of the data will determine whether discretionary access controls (DAC) or mandatory access controls (MAC) are required. DAC allow users to set access permissions on their own files as they wish (e.g. the standard Unix file permissions). MAC are set and enforced by the AIS operating system security kernel based on security labels assigned to system objects (e.g. files, directories, processes and users) in accordance with the organization's security policy. In the latter case, users do not have complete control over how the permissions are set. MAC is used typically in multilevel security environments. (Russell and Gangemi, 1991; DoD, 1985) The NPS ADP Security Program directive (NPS, 1992) does not address general data confidentiality and integrity policies for unclassified computing systems.

3. System Backups

System hardware is usually replaced rather easily, although often at considerable one-time cost. The same cannot always be said about the data that is stored on a system; data is often priceless or is not reproducible. Data and software are vulnerable to a variety of threats that include viruses or other malicious software, disgruntled or untrained employees, and hardware failure (e.g. a hard disk crash). The easiest control against unintentional or malicious modification or destruction of data is to conduct adequate system backups. System backups must also be stored securely to ensure they themselves are not damaged or otherwise corrupted in the event of an incident. Backups must be a true representation of the system state before the incident. (Fraser, 1996; NPS, 1992; Russell and Gangemi, 1991; Wong and Watt, 1990)

Again, the importance of the data stored on the system will determine the frequency of the backups. The organization will want to backup critical data much more often than noncritical data, but even noncritical data needs to be backed up periodically. Although it is wise for users in a multi-user environment to maintain personal backup copies of their files, this is often infeasible due to classification or large file sizes. Therefore it is appropriate for users to expect the computing services staff to conduct regular file system backups that include even noncritical user data. The organization's backup policy must consider user expectations and must also be documented and widely distributed within the organization. NPS backup policy is documented in (NPS, 1992) and explicitly states that users are responsible for their own data files. Individual system backup requirements vary depending on the classification of information that is processed or stored on each individual system (NPS, 1992). While there is some disagreement between the users and administrators concerning the adequacy of the NPS backup policies, the policies are widely followed.

E. TELECOMMUNICATIONS AND SYSTEM ACCESS ISSUES

Managers must consider telecommunications and system access issues when determining system security requirements. Subsection 1 discusses network access controls. Subsection 2 discusses the concepts of identification, authentication and authorization. Subsection 3 discusses the notion of nonrepudiation. Subsection 4 discusses data confidentiality and integrity with regards to network communications. Subsection 5 discusses issues relating to network connectivity.

1. Access Control

Network access security is arguably the most important and most challenging facet of AIS security given today's computing environment. Every unprotected network port is an open door into the system. As with system availability requirements, network access security is multifaceted. It concerns not only physical access to system components (i.e. the network infrastructure) and access to data and software as discussed above, but also telecommunications access. Intruders may gain access to a network through telecommunications networks to manipulate the data or software. NPS access control policy is to grant access to only those systems and data an individual needs to accomplish assigned tasks. This policy is specified in (NPS, 1992). The general policy is widely followed but specific procedures for implementing the policy are left to the discretion of individual network managers.

The first line of defense for network access security is controlling physical access to system hardware components (i.e. the network infrastructure) including workstations, network cabling, disk drives, communication routers and network switches. Although controlling access to private network components through routine physical security mechanisms is relatively straightforward, controlling access to the public infrastructure is not possible. Public routers and switches can be subverted; an intruder can eavesdrop on network traffic by tapping cables and wires or by gaining access to network switches and routers and reconfiguring the routing tables to misroute data packets. Public networks are outside the realm of responsibility of the organization and must not be trusted. (Kluepfel, 1996)

A physically secure network does not guarantee network security. Modern computer networks are typically protected from unauthorized telecommunications access by a filtering-router firewall. Filtering-router firewalls examine the source IP address of incoming data packets and filter out those packets with IP addresses that are not authorized access to the network. Unfortunately the source IP address of incoming data packets can be altered by an intruder to fool a firewall into allowing otherwise unauthorized packets access to the network. This type of attack is called "address-spoofing." A number of spoofing methods are documented in (Bellovin, 1989). Once the firewall has been breached the network applications that use source IP addresses for authentication are vulnerable to the same form of attack. (CERT, 1995; Bellovin, 1989)

The "cracker" (or "hacker")[7] threat has been widely publicized in recent years. In 1996, unclassified DoD computer systems were attacked an estimated 250,000 times with nearly two-thirds of those attacks being successful (Zuckermann, 1996). "Cracker" attacks include attacks on the network infrastructure as well as attacks against the software and data. These attacks include password-sniffing attacks, address- spoofing attacks and denial of service attacks. Most successful attacks exploit well-known, well-documented system vulnerabilities. Therefore it is imperative that system administrators remain informed about and attempt to remedy system vulnerabilities. Of course it is important to note that system administration personnel must expend valuable time and resources to remain current. Like most security issues, there is a tradeoff. (Kluepfel, 1996; Haller, 1994; Rich, 1992)

2. Identification, Authentication and Authorization

The next line of defense in network access security is identification, authentication and authorization. Identification verifies the existence of a particular user. Authentication verifies the user is actually who they claim to be. Authorization verifies the user is allowed access to particular system services. These three network security functions form the foundation of all AIS security; an intruder cannot gain access to an AIS unless they subvert all three functions. Controlling access to the public network

[7] The terms "hacker" and "cracker" are somewhat ambiguous and greatly debated. (Russell and Gangemi, 1991) defines a "hacker" as an individual with a keen interest in computer systems and an eagerness to experiment with them and test their limits. They define a "cracker" as an individual who intentionally breaks into computer systems with malicious intent.

infrastructure is the ideal first step in network access security that is not feasible. Therefore identification, authentication and authorization become the first line.

The first objective of most attacks on a system is to gain access to the data. To successfully gain access an attacker need only determine the user ID and password for a single user on the system. Unfortunately users are notorious for choosing poor passwords or are careless in protecting them. Conversely, "good" passwords are difficult to remember which introduces additional insecurities when users write down their passwords or are unable to access the system when they forget their passwords. Common techniques for obtaining passwords include "social engineering" (e.g. impersonating a system administrator over the telephone) (Winkler, 1995), exploiting operating system vulnerabilities to gain access to the system password file, installing a Trojan Horse[8] on a networked workstation to transmit login sequences to the intruder, and capturing unencrypted remote login sequences by "sniffing" data packets while eavesdropping on remote network connections. One of the most significant dangers associated with static multi-use passwords is that it usually is not possible to detect or know when a particular password has been compromised.

In general, using static passwords has been greatly criticized by law enforcement, academic and government experts. These experts recommend implementing a dynamic password scheme such as Bellcore's S-Key protocol (Haller, 1995) or other secure access

[8] A Trojan Horse is a computer program or code fragment that "hides in an independent program that performs a useful or appealing function — or appears to perform that function. Along with the apparent function, however, the program also performs some other unauthorized operation" (Russell and Gangemi, 1991, pg. 83).

mechanism such as MIT's Kerberos protocol (Miller *et al.*, 1985). Implementing a second authentication mechanism such as a magnetic smart card or a biometric device such as a fingerprint reader will greatly increase system and network access security. (McNulty, 1994; NIST, 1994; Bennington, 1991)

3. Nonrepudiation

Nonrepudiation is the process of protecting against denial of receipt or transmission of data by the receiver or sender respectively. As the commercial use of the Internet continues to expand, nonrepudiation will become increasingly important. Electronic commerce will require nonrepudiation mechanisms to protect the interests of both suppliers and consumers. This will require some form of cryptography. Many current nonrepudiation mechanisms are based on the Digital Signature Standard (DSS) (NIST, 1992). There are many issues and concerns that still must be resolved before nonrepudiation mechanisms are widely implemented. The greatest of these include whether to use secret-key cryptography or public-key cryptography, how to distribute and manage the cryptographic keys, what control an organization will exercise over its employees' keys, what role government will play in the process and what impact cryptographic software export restrictions will have on the organization. These issues will need to be addressed in the organization's information security policy. (Carpenter and Baker, 1996; Fahn, 1993) Nonrepudiation policies are currently uncommon due to lack of consensus on technical, administrative and economic issues.

4. Integrity, Confidentiality and Authenticity

Because securing the network "pipe" is infeasible in today's interconnected networking environment, we must secure the data in the "pipe" if data security is required. Confidentiality during network transmission can be sufficiently assured by using a bulk encryption protocol such as DES (NIST, 1988) or RSA (Rivest *et al.*, 1978). If encryption is done below the network layer, intermediate nodes within the network must be able to decrypt the data packet to find where it next needs to be routed. This slows delivery and also provides the opportunity for an intruder with access to an intermediate node to view the packet contents. Encrypting the data packets at the application layer guards against this possibility, but also opens the door for intruders to analyze traffic flow between nodes (i.e. packet frequency, size, addressees etc.) with the hope of piecing together useful information. Just as with most decisions concerning AIS security, there are many tradeoffs.

Data integrity during network transmissions can be sufficiently assured using cryptographic checksums or message digests such as MD5 (Rivest, 1992) or the secure hash algorithm (SHA) portion of the DSS (NIST, 1992). Similarly, data authenticity can be sufficiently assured using digital signatures such as those provided by the DSS or PGP (Zimmerman, 1994).

5. Network Connectivity

Besides the physical security of network cabling and hardware, there are three connectivity variables in which we are interested. The first variable is the type of hosts to which we are connecting. The security challenges present when connecting similar hosts (e.g. ATM-to-ATM or IP-to-IP) may be different than when connecting dissimilar hosts (e.g. ATM-to-IP). The second variable is the connection method. The security challenges when remotely connecting to a network over the Internet may be different than when remotely connecting to a network using direct-dial telephone lines. The third variable is the message address. Messages with multicast addresses may be vulnerable to different security challenges than messages with unicast addresses. We examine these network connectivity issues in detail.

a. Similar versus Dissimilar Hosts

Much current work focuses on running IP and other higher-layer protocols over ATM in addition to native ATM-to-ATM connections. This hybrid approach theoretically permits the use of TCP/IP firewalls to protect a network. However, native cell-based ATM environments cannot use IP packet filters because protocol differences are too great (Rendleman, 1995). Nevertheless research is being performed in the area of ATM firewalls (Chuang, 1995; Raynovich, 1995). One approach recommends filtering only the call setup cells (Chuang, 1995). Another proposed solution combines hardware and software and promises to filter each cell without slowing down the data rate (Raynovich, 1995). In either case, for hosts connected to both an internal Ethernet LAN and an external ATM WAN, there is little assurance (other than protocol complexity, a dubious safeguard) that an intruder will not come through an ATM port, exploit a system vulnerability and then gain access to the internal IP network (Rendleman, 1995).

b. Internet versus Direct-dial Telephone Lines

The Internet is an open environment and various insecurities have been discussed throughout this work. As mentioned above, most Internet attacks attempt to exploit well-known, well-documented system vulnerabilities. While patches are available for most of these vulnerabilities, it is essentially impossible that all nodes have the latest patch; there will always be a hole somewhere in the net. It is unwise for organizations to trust or expect Internet Service Providers (ISPs) to provide security for data transmitted over the Internet regardless of what the underlying network infrastructure looks like. Since the Internet is not an inherently secure environment, it is similarly unlikely that a public ATM network will be a particularly secure environment.

Direct-dial telephone lines are safer but are still vulnerable. Communication networks are weakly protected and particularly susceptible to human interference. The following quote from an IETF draft copy of *National Information Infrastructure Risk Assessment: A Nation at Risk* as quoted in (Bollentin, 1996) demonstrates the vulnerability quite well.

> Within the PSN [public switched network], intruders have already compromised nearly all categories of activities, from switching systems to operations, administration, maintenance, and provisioning (OAM&P) systems, and to packet data networks. Private branch exchanges and corporate networks that tie into the public network have crashed or disrupted signal transfer points (STPs), traffic switches, OAM&P systems, and other network elements. They have planted destructive 'time bomb' programs designed to shut down switching hubs, disrupted E-911 services throughout the eastern seaboard, and boasted that they have the capability to bring down all switches in Manhattan. [Bollentin, 1996, pg. 20]

c. *Unicast versus Multicast Message Addresses*

Most Internet communications rely almost exclusively on unicast (i.e. point-to-point) messages. Messages from one host typically are addressed and sent to only one other host. Even multi-address electronic mail messages are transmitted as unicast IP datagrams since the e-mail message is replicated at the mail server and a copy is sent to each addressee. Internet security (and hence that of unicast messages) have been discussed throughout this thesis and is documented in a variety of sources (Fraser, 1996; Klaus, 1995; Ranum, 1995; Curry, 1990).

Despite its predominantly unicast nature, the Internet also supports multicast (i.e. multipoint-to-multipoint) messages using the Internet Multicast Backbone (MBone). The MBone is a part of the Internet and therefore exhibits the same vulnerabilities. As with the unicast Internet, security is provided primarily by individual MBone applications. Most of these applications pose little threat to end systems as they cannot be "coaxed" into writing to disk by incoming packets nor do they run *set-uid*.[9] However, MBone management tools such as *mtrace* or *mrinfo* do run *set-uid* and other MBone management tools such as *map-mbone* require root privileges. (Kumar, 1996; Savetz, *et al.*, 1996) Giving public access to these tools, either directly or via a script imbedded in a home page, introduces a possible system vulnerability. The likelihood an automated script can be used to invoke these tools fast enough to disrupt network service or that

[9] *set-uid* is a flag used in remote procedure calls that when set for the called procedure, changes the calling process' *uid* to that of the called procedure. If the called procedure is running as *root*, then the calling process will also run as *root* with all the associated privileges until the called procedure returns (Curry, 1990).

these applications can be subverted by an intruder to gain access to system files is uncertain but the vulnerability exists nonetheless. A recent software analysis of these monitoring applications indicates that they are adequately secure for open use (Erdogan, 1996).

The MBone is also theoretically vulnerable to Trojan Horses disguised as MBone applications. For example, someone might advertise that a new interactive whiteboard (wb) application which fixes bugs in previous versions is available for download from the Internet. This seems to be a useful application and one in which people might be interested. Imagine the frustration if embedded within this appealing application was the command *rm -rf.* When such a new interactive "whiteboard application" is executed all files below the current directory will be quietly erased. Currently the Trojan Horse threat is of little concern as there are relatively few MBone applications. These applications are written by a relatively small number of individuals who are trusted and closely monitored within the MBone community. This is an example of consensual administrative policies providing group security. But as MBone use increases and the multicasting influence expands, the variety and number of MBone applications is likely to increase thus increasing the threat.

The MBone is also vulnerable to threats relating to security firewall configurations. The first vulnerability is the susceptibility to multicast address spoofing. Multicast messages are sent over the regular Internet but require special routers to handle the multicast addresses. As multicast routers are not widespread, multicast packets are encapsulated within standard unicast packets. Tunnels are established between multicast routers through the mesh of regular unicast routers. The vulnerability exits not so much in the MBone itself, as in network security firewalls. Most security firewalls examine host IP addresses (in this case the mrouter addresses at the tunnel endpoints) accepting packets from authorized hosts and rejecting packets from unauthorized hosts. However, packets traversing the multicast tunnel are not authenticated, and therefore any IP host can send malicious packets through the firewall by posing as the remote tunnel endpoint. (Kumar, 1996)

Directly related to the previous vulnerability is the fact that current MBone applications use only the User Datagram Protocol (UDP) at the transport layer, but multicast addresses are at the IP layer. Because most firewalls filter on the host IP address (e.g. the IP address of the mrouter serving as the remote tunnel endpoint), all UDP packets on subscribed multicast channel/port combinations are received by the operating system kernel regardless of the UDP packet contents. Users may think they are receiving audio packets when in actuality they are receiving a malicious network file system packet. To protect against such an attack, filtering routers must inspect the UDP payload of incoming IP packets for any unwanted UDP ports or protocols. Trusted software must also behave robustly in the presence of unexpected inputs. (Savetz *et al.*, 1996)

As with all network security decisions, security for the MBone is a tradeoff. MBone applications can be reasonably trusted at present because most applications are written by "trusted" individuals. This may change as multicasting becomes more widespread. Multicast packets can also be filtered by firewalls for added security.

However, a balance is required between the two filtering extremes. Filtering granularity that is too fine will degrade performance and may unnecessarily restrict access to MBone services. Filtering granularity that is too coarse may leave the internal LAN vulnerable to attack. To date, malevolent security incidents on the MBone have been rare or nonexistent.

Other threats associated with the MBone also have been identified, most affecting system availability or the quality of service (QoS). For example, one user leaving a microphone on or transmitting high-bandwidth video can temporarily degrade network services for every other user (Macedonia and Brutzman, 1994). Setting the time-to-live variable too high may also flood the MBone and make it unusable. As with any other secure channel, trusted distribution of cryptographic keys for "private" MBone sessions also remains a separate issue. In general, administrative policies (such as Internet Service Providers detecting/isolating offenders and community cooperation) have been remarkably effective in minimizing system availability and QoS problems. Lessons learned are typically automated and incorporated into global software distributions.

F. IMPLEMENTATION ISSUES

1. Effectiveness and Efficiency

Because AIS security competes for scarce organizational resources (i.e. labor hours and capital), the resources that are allocated to AIS security must be justified. Controls must be both effective and efficient. A control that does not adequately protect against the threat for which it was implemented is not effective. A control that protects against the threat but is too costly to implement, operate or maintain (compared to the value of the hardware, software and data it is protecting) is inefficient. Managers must carefully consider the cost/benefit relationship between AIS security controls and the value of the information and equipment the controls are protecting given the existing environment. NPS policies regarding effectiveness and efficiency are currently decided on a LAN by LAN (typically department by department) basis.

2. Life-Cycle Costs

To adequately justify security controls, management must understand the costs associated with AIS security. AIS security increases the cost of an information system over the entire lifespan of the system. Costs are incurred to implement controls. Costs are incurred to operate, maintain, support, monitor and enforce the security policy and the controls that are implemented. All controls interfere with productivity in one fashion or another and thereby also increase system costs — hence the less obtrusive the controls, the lower the cost incurred. Finally, AIS security controls reduce a system's life span by increasing a system's complexity and reducing its flexibility — the more complex and less flexible a system, the sooner the system will need to be replaced. This requires the organization to expend resources sooner than it otherwise might. (Baskerville, 1988)

NPS life-cycle planning is performed on a component by component basis. NPS life-cycle planning requirements are briefly described in (NPS, 1992).

3. Convenience and Acceptance

Managers do not exist in a vacuum; other stakeholders — administrators, technicians and users — are affected by their decisions. This often complicates the AIS security problem due to occasional lack of consensus among the various system stakeholders on just how vulnerable the system is and what needs to be done about it. It does not much matter what security controls are implemented if such controls unduly inconvenience personnel. Security controls that are excessively complex or restrictive will soon be circumvented by IS personnel and users thereby rendering such controls useless (Fitzgerald, 1993). Extended system shutdowns for security repairs (as in the December 1995 NPS incident) jeopardizes daily operations, mission completion and long-term sustainability.

Ideally, security controls are designed into a system from its inception, taking user and mission needs into account. This practice tends to reduce the productivity costs associated with AIS security by making control and audit mechanisms transparent to system users and reducing user resistance to such mechanisms. However, security requirements are often overlooked during the requirements-analysis phase of system design. (Fraser, 1996; NCSC, 1993; Baskerville, 1989; Palmer and Potter, 1989) Similarly, if IS administrators, technicians and users do not understand the need for system security, any control that is implemented will likely be circumvented. It is therefore paramount that IS personnel be made aware of IS security needs and priorities through training and continued communication between and *among* managers, administrators, and users. (Fitzgerald, 1993; Wong and Watt, 1990)

G. SUMMARY

Securing a networked computer system is a complex problem. If a networked computer system is to be secure, managers must understand the need for system security and decide how much and what type of security controls are required. Their decisions and priorities must be well documented and widely distributed in the form of an organizational security policy. The security policy must clearly state the rules and goals that the system security controls will enforce. Managers need to ensure adequate resources are allocated to effectively implement, monitor and enforce the security policy.

Managers need to understand the environment in which their organization's computer system/network is operating and the threats which the environment presents. By understanding the enemy, so to speak, managers can develop and implement controls, plans and procedures to either prevent a security incident from occurring or to effectively deal with it if one actually occurs. By being proactive rather than reactive, the organization may be able to recover from an incident much more quickly and at far less cost than it otherwise might.

The foundation of networked computer system security is in controlling both physical and telecommunications access to the system. Telecommunication access control is accomplished through identification, authentication and authorization mechanisms supported by cryptographic software. Access control becomes more complex when access is allowed over both the Internet and direct-dial telephone lines. Nonrepudiation, data integrity and data confidentiality are becoming more important as organizations begin to use networking and the Internet for commercial gain. Nonrepudiation, data confidentiality and data integrity also can be accomplished with digital signatures and cryptographic algorithms. Finally, while not necessarily unique, multicasting does include certain vulnerabilities that have been successfully prevented to date through trusted software and effective administrative policies.

Security controls must be carefully selected based on their effectiveness, ease of use and efficiency. Complex controls that prove too restrictive or difficult to use will simply be circumvented by exasperated users and administrators. Of course, all system users must be made aware of and understand the need for AIS security. Increased awareness and understanding can be achieved through training and frequent communication between and among managers, AIS personnel and users. Time and resources spent on training can pay dividends by increasing the likelihood that security controls are understood and followed.

V. SOFTWARE INSTALLATION: KERBEROS

A. INTRODUCTION

Section B in this chapter explains the motivation for implementing the Kerberos authentication and authorization protocol. Section C presents a general discussion of the Kerberos protocol and how it works. A detailed analysis is not presented here as more in-depth discussions of the protocol can be found in (Kohl *et al.*, 1992), (Steiner *et al.*, 1988) and (Miller *et al.*, 1987). Section D discusses some of the planning issues that must be considered before any meaningful implementation of Kerberos is undertaken.

B. MOTIVATION: PASSWORD-SNIFFING PROTECTION

One of the most common attacks against network security is a packet-sniffing or eavesdropping attack during remote sessions. With standard Unix remote access applications (e.g. *rlogin*, *rsh*, *rcp*, *ftp*, *telnet* etc.), passwords are transmitted from the remote system to the host system over the network in the clear (i.e. unencrypted). An eavesdropper at any intermediate node between the remote system and the host can easily capture the login sequence (of user ID and password). The captured sequence can then be used for future unauthorized access to the host system — a particularly undesirable consequence of remote computing. This type of attack has been rather successful in the past. For example, in February 1994, the Computer Emergency Response Team (CERT) at the Software Engineering Institute (SEI) at Carnegie Mellon University issued an advisory for all Internet users to change their passwords in the wake of numerous successful packet-sniffing attacks (CERT, 1994).

Academic users often need to travel to conferences and maintain Internet connectivity to securely collaborate over the Internet with colleagues at remote locations. Thus password security is an important requirement. One solution to the problem of transmitting passwords over an insecure network is the Kerberos authentication and authorization protocol.

C. THE KERBEROS PROTOCOL

This section briefly explains the Kerberos protocol. Subsection 1 provides a short description of Kerberos and provides some background information with regard to the protocol's history and development. Subsection 2 provides a general overview of how the authentication protocol works.

1. Background

Kerberos is a client/server-based access security protocol that relies on a trusted host known as the "Kerberos authentication server," the "Kerberos key distribution center (KDC)," or simply "Kerberos." Kerberos provides authentication services for network principals that are acceptable to both the client and the server. A network principal is any user (or a client process acting on behalf of a user) or service on a network. The trusted host (i.e. “Kerberos” host or simply the “KDC”) stores a secret password for all principals on the network. Any principal wishing to access the network or a protected network service must first authenticate itself to the KDC before access is granted.

Kerberos was developed in 1987 by researchers at the Massachusetts Institute of Technology (MIT) to protect emerging network services provided by the school's *Project Athena*. When Kerberos was designed, MIT was operating 750 computers on 30 subnets and supporting over 5000 active users. MIT's computing environment consisted of both distributed untrusted workstations and time-share mainframe-based systems, but there was virtually no centralized control of the computing resources on the MIT campus and the individual workstations were unsecured. Malicious users might subvert or otherwise gain control of individual workstation operating systems and gain unauthorized access to the network or masquerade as another user to gain unauthorized access to files. (Miller *et al.*, 1987)

The primary purpose of Kerberos was to extend the familiar time-sharing notions of authentication, authorization and auditing to the distributed network environment operating with untrusted workstations. To that end, Kerberos had three goals: (1) provide both one-way and mutual authentication between principals to the granularity of a single user and a single network service instance; (2) provide coarse access authorization and allow services to implement specific authorization models as required while assuming reliable authentication of users; and (3) allow integration of a modular accounting system to provide accounting services for resources as required (Miller *et al.*, 1987). While the Kerberos protocol does have some limitations and weaknesses, the protocol effectively achieves these goals (Bellovin and Merritt, 1990).

2. The Kerberos Protocol

a. Authentication

Kerberos is a client/server-based protocol that relies on a trusted third party to provide authentication services for both a network client and the network service the client wishes to access. Thus any network application can be customized to comply with the Kerberos authentication protocol. The Kerberos protocol does not transmit passwords *per se* over the untrusted network, but rather authenticates users and processes at the client machine. The protocol uses a modularized 64-bit encryption algorithm to accomplish the authentication function. The default encryption algorithm is the Digital Encryption Standard (DES) (NIST, 1988), but any 64-bit encryption algorithm can be used. The protocol is flexible in that the same general principle is used both for the initial network login sequence and for subsequent network application and service access. Figure 6 illustrates the general authentication protocol.

When initially accessing the network a client must request authentication from the centralized Kerberos key distribution center (KDC) (i.e. Kerberos). Integral to the KDC is a database that stores the password for every user and each "kerberized" network application or service. Upon verifying the ID of the user, the KDC issues the client an initial encrypted authentication ticket called a “ticket-granting ticket” (TGT). This ticket is encrypted using the user's stored password as the seed for the encryption algorithm. To decrypt the TGT, the client process prompts the user for their password. The user will only be granted access to the network if they enter the correct password to decrypt the TGT.

At this point the user is only logged onto the network; they can access standard applications or services but they cannot yet access Kerberized applications or services. To make the protocol as transparent to the user as possible, the protocol requires the user to enter their password only once during any one session. The password

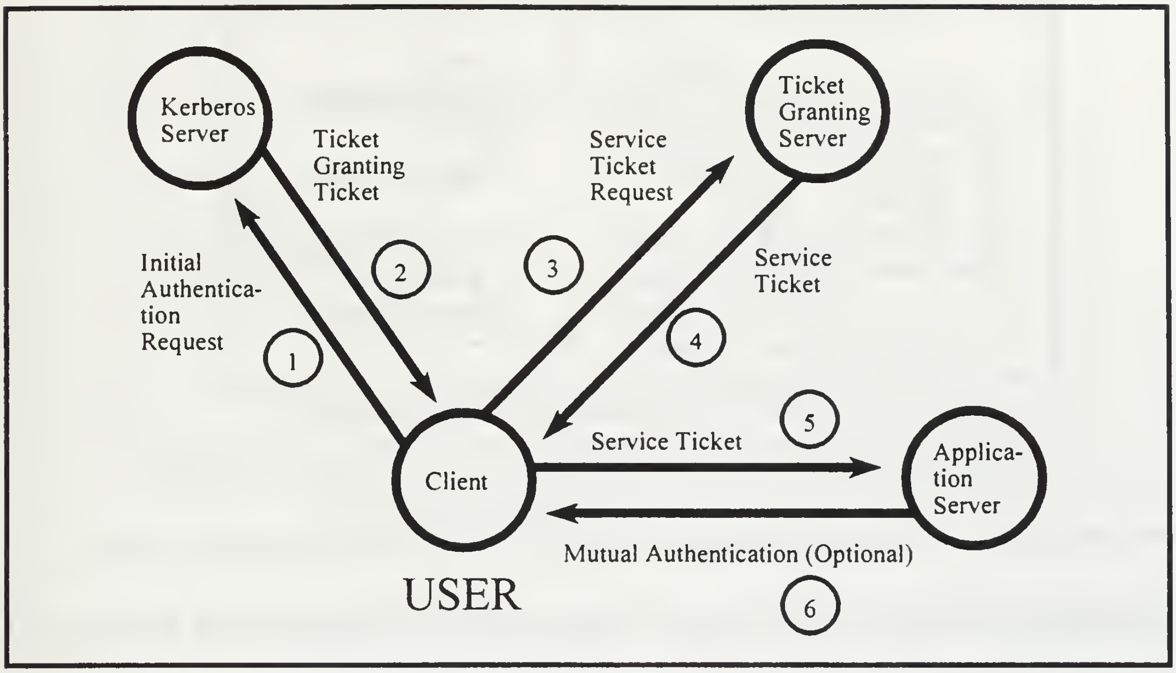

Figure 6. The Kerberos authentication protocol. After (Steiner, *et al.*, 1988)

is kept in memory on the client machine only long enough to decrypt the TGT. The TGT is then used in lieu of a password to obtain authentication tickets for individual network applications and services as required by the user.

Once logged onto the network, a user wishing to access a network application or service must obtain an authentication ticket from a ticket granting server (TGS) by submitting an authentication request and the TGT to the TGS. The TGS may or may not reside on the same machine as the KDC. Once obtained, the application or service ticket is then presented to the particular application server to gain access to the application or service.

The structure of each ticket (i.e. the TGT obtained from the KDC and each specific service ticket obtained from the TGS) is identical and is illustrated in Figure 7. Inside each ticket is a session key for use between the client and the server portion of the

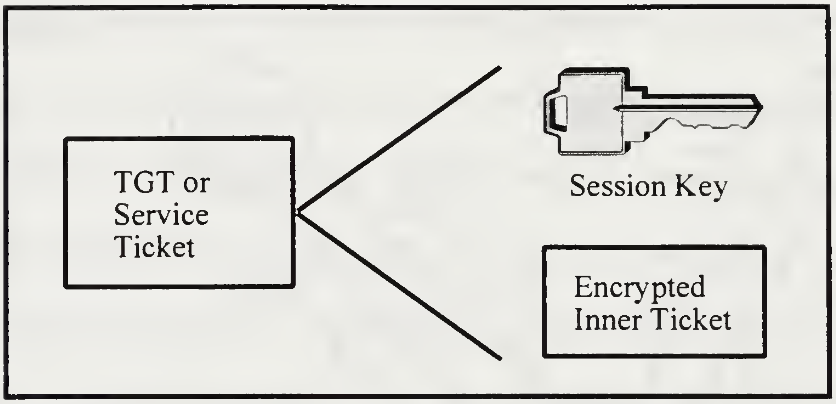

Figure 7. Contents of a Kerberos authentication ticket.

network application or service (including the TGS in the case of the initial TGT). Each authentication ticket issued to a client process also contains another ticket that is encrypted with the application's (or TGS's) password. This inner ticket is not readable to the client (or the user on whose behalf the client is operating); it may only be decrypted by the server portion of the applicable application or service.

The inner ticket has its own structure as shown in Figure 8. Inside the inner ticket is the client's name, the client's IP address, the server's name, the current computer time, the ticket's lifespan, and the same session key that was issued to the client as part of the authentication ticket, whether the authentication ticket is the original TGT or is a specific service ticket.

To access a particular "kerberized" network application the client process will prepare an authenticator composed of the client's name, IP address and the current time. The authenticator is encrypted with the session key issued with the initial TGT. The client will present this encrypted authenticator, the name of the specific service and the encrypted inner ticket (i.e. the actual TGT) to the ticket granting server (TGS). The

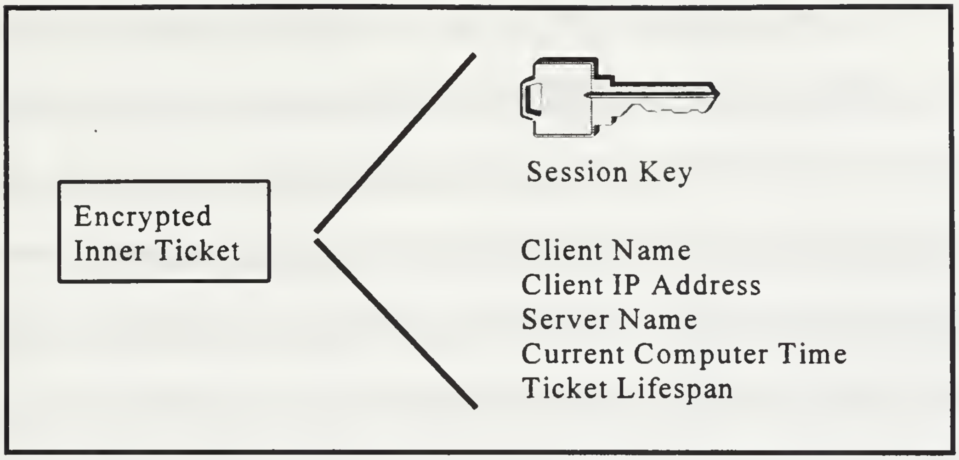

Figure 8. Contents of the encrypted inner ticket within a Kerberos authentication ticket.

TGS decrypts the encrypted TGT with its unique “password”and retrieves the session key from within the TGT. The TGT then decrypts the client's authenticator with the session key and compares the information contained in the TGT to that contained in the authenticator. If the client name and IP address match, and the time stamp on the authenticator is not beyond the lifespan of the TGT, the TGS authenticates the client and issues an authentication ticket for the specific application service specified in the request. This new ticket has the same structure as the original TGT but a different session key for use between the client and the application server.

The authentication process to the specific application server is identical to that which has just been described. The authenticator and the encrypted inner ticket are now presented to the specific application server.

b. Authorization

Kerberos provides the flexibility to allow different network services to implement different authorization models (access control lists for example). There are three ways the Kerberos authentication protocol can be extended to include authorization. The first is by including authorization information for each network service directly in the Kerberos database. Only authorized users are then granted specific service tickets. The second means of providing authorization is for the system to maintain separate access control lists (ACLs) for each service. The Kerberos protocol can be modified to provide certified copies of the ACL in place of an authentication ticket. The certified ACL is then presented to the ultimate service in lieu of the authentication ticket. The third way the Kerberos protocol can be modified to provide authorization is by customizing each individual service to maintain its own authorization information. The third authorization method is recommended by the Kerberos developers at MIT. Library modules implementing this authorization model using ACLs are included in the standard MIT distribution of Kerberos. (Miller *et al.*, 1987)

While the Kerberos authentication model is centralized, the authorization model is decentralized. The principle idea behind decentralized authorization is that each service is likely to know best who the authorized users are and how to best implement that authorization. The advantages of this model are outlined in (Miller *et al.*, 1987) and will not be addressed here.

c. Accounting

Kerberos does not implement an accounting system itself, but it can accommodate a modular accounting system to track resource usage (Miller *et al.*, 1987). The modularized accounting system is used for tracking system resource usage. That information is usually used for an organization's charge-back/computer resource billing policy. If a network application (e.g. the print daemon) is “kerberized” and the organization implements a charge-back policy, then Kerberos can accommodate a modularized accounting system to gather the necessary data.

For discussion’s sake, it does not seem feasible to use a modularized accounting system to evaluate usage and practicality of the Kerberos system itself as Kerberos does only one thing: control access to system resources. Kerberos accomplishes this function through authentication and authorization. The Kerberos system itself is Boolean in nature: either you use it or you do not; either a network application is “Kerberized” or it is not. If an organization implements Kerberos to protect network resources, then by default Kerberos gets used.

Evaluating the practicality of Kerberos is equally problematic. As with most computer system/network security mechanisms, evaluating the practicality of Kerberos is similar to precisely evaluating the practicality of life insurance. How does one do that? The nice thing about Kerberos is that it can replace the standard login sequence and be transparent to the user. Therefore it is only as obtrusive as the regular login sequence. This is assuming of course that network managers and system administrators want to implement Kerberos and take advantage of the access protection it provides.

D. PLANNING REQUIREMENTS

Specific planning requirements will vary from organization to organization. This section simply highlights some of the issues that must be addressed before implementing Kerberos within an organization. Subsection 1 addresses network organization issues.

Subsection 2 discusses network service issues. Subsection 3 addresses network administration issues.

1. Network Organization and Realms

Any large-scale implementation of Kerberos (such as those at MIT (Miller, *et al.*, 1987)) requires considerable planning. Kerberos operates within the concept of a *realm* which is similar to an Internet domain. However a realm is fundamentally different from an Internet domain. A domain is based on a common set of network addresses that can be divided easily into subdomains by simply parsing up the common set of addresses into subsets which can be allocated to form subnetworks. On the other hand, a realm is an independent administrative division or grouping of computers that cannot be divided into "subrealms." As an application layer protocol, Kerberos uses lower level protocols only for transmission and delivery. A host's network address is not a prerequisite for inclusion in a particular realm. (Miller, *et al.*, 1987; Neuman and Steiner, 1988)

Network-administrators deploying Kerberos must decide how many realms they are going to operate at their site and how these realms will be organized. To simplify network management it is recommended that realms mirror an organization's top-level Internet domain name, however this may be impractical or inefficient. (Neuman and Steiner, 1988). Small sites can operate adequately within a single realm; larger sites will require multiple realms. For example, MIT operates with three realms supporting more than 5,000 users and 750 computers on 30 subnets (Neuman and Steiner, 1988). These decision outcomes will depend largely on the size of the site and the number of users at the site.

2. Network Services

Managers must also decide how many application servers will be installed within each realm, where those application servers will be located both physically and logically at the site, and what network applications/services will be protected. This section discusses these decision issues.

a. Application Servers

(Eichin and McGregor, 1995) recommends making every local host an application server so users can access them directly. While this may be desirable, it may be impractical for large sites. Because a single master KDC has complete authentication and administration authority within a particular realm, making every local host an application server at a large single-realm site will greatly increase the workload on the single master KDC. This will also increase the size of the KDC's database and make database management more difficult.

The server workload issue can be solved by implementing slave servers that contain a read-only copy of the database, but the database management issue remains. Both the server workload and database management issues can be solved by subdividing a large realm into a collection of smaller realms. In this case authentication between realms is desirable but is not fully supported by the current version of Kerberos (version 4) (Miller, *et al.*, 1987). Cross-realm authentication is supported with version 5, but version 5 is still in the beta testing phase with many bugs left to iron out (Kohl and Neuman, 1993).

The decision concerning which applications and services to protect using Kerberos is closely related to the earlier decision concerning which hosts will be application servers. It may be desirable to have every network host provide some basic set of "kerberized" services (e.g. *rlogin*, *telnet*, *rsh*, *rcp*, *ftp* etc.), but it is also reasonable to limit some applications (e.g. printer daemons and World Wide Web servers) to a few specific machines. Freeware distributions of Kerberos include "kerberized" versions of the standard Unix r-commands (i.e. *rlogin*, *rcp* and *rshell*) and most freeware distributions also include "kerberized" *telnet*, *tftp*, *ftp* and *pop* (e-mail server) applications. Freeware distributions of Kerberos also include a library of procedures and subroutines programmers can use to "kerberize" existing applications. The time and expense of converting applications will depend on programmer skill and the complexity and size of the applications managers wish to "kerberize" and make available to users. This remains an area for further research since thesis does not address specific costs.

Another closely related issue concerns where within an organization application servers will be placed (both physically and logically with respect to network addressing). This point is moot for those sites that make every host an application server for every application. Since in most sites it is unlikely every host will support every application, this issue must also be addressed.

b. Clients

Every network host (including the application servers) for a legitimate implementation of Kerberos will be a Kerberos client. The decision to limit the accessibility of other "kerberized" applications to specific hosts is no different than under standard non-Kerberos practices.

3. Network Administration

Network administration when operating with Kerberos is only slightly more complex than when operating without Kerberos. The added complexity results from having to manage an additional database containing passwords for both network users and for the network services. However, depending on an organization's security policy, the standard passwords can be disabled thereby eliminating the need to separately manage standard passwords. The motivation here is to make the Kerberos initialization process transparent to the user by replacing the standard login script with one that handles the Kerberos initialization. When a user logs into a workstation they are automatically initialized to the KDC for that realm. All authentication requirements within that realm will then be handled by Kerberos. From a user's standpoint the whole process is invisible. In either case user accounts must be managed, users and network administrators alike must be trained, network services must be managed, network access must be monitored, etc. A well-planned implementation can eliminate many administrative burdens.

E. SUMMARY

The Kerberos protocol provides authentication of users and client processes by means of authentication tickets. The protocol also supports authorization and can accommodate modular accounting systems to track system usage. This chapter explains the motivation for implementing the Kerberos protocol and briefly describes how the protocol works. This chapter also discusses many of the planning issues that must be addressed prior to an operational implementation of Kerberos. Kerberos is of interest to NPS as a means to provide remote password protection to seamlessly and securely extend LAN connectivity across the Internet.

VI. INTEGRATED IP/ATM LAN/WAN SITE SECURITY SURVEY

A. INTRODUCTION

The IP/ATM LAN/WAN is located in the System Technology Lab (STL) at NPS. The STL contains both classified and unclassified computer networks/systems, but this thesis addresses only the unclassified systems. For simplicity, the abbreviation STL is used throughout this chapter to designate the unclassified system. This chapter documents the results of a site security survey conducted for the unclassified STL IP/ATM LAN/WAN. This chapter is not a risk analysis *per se*, but merely a careful look at the level of security for the unclassified networked computer system as well as what work still needs to be done. Conclusions in this chapter are drawn from a comparison of the information obtained in background research, structured interviews and free-form interviews with the STL computing staff. Issues are described in detail in Chapter IV. Section B briefly describes and illustrates the integrated IP/ATM LAN/WAN network topology. Results of the STL site security survey are presented in section C.

B. INTEGRATED IP/ATM LAN/WAN TOPOLOGY

Figure 9 illustrates the topology of the integrated IP/ATM LAN/WAN at NPS. (Courtney, 1996) explains this topology in detail. In general, the core of the network consists of two Unix workstations connected together through both an IP Ethernet interface and an ATM interface. On the left side of Figure 9, the IP Ethernet interfaces connect the workstations to the STL subnet, which is connected to the campus backbone

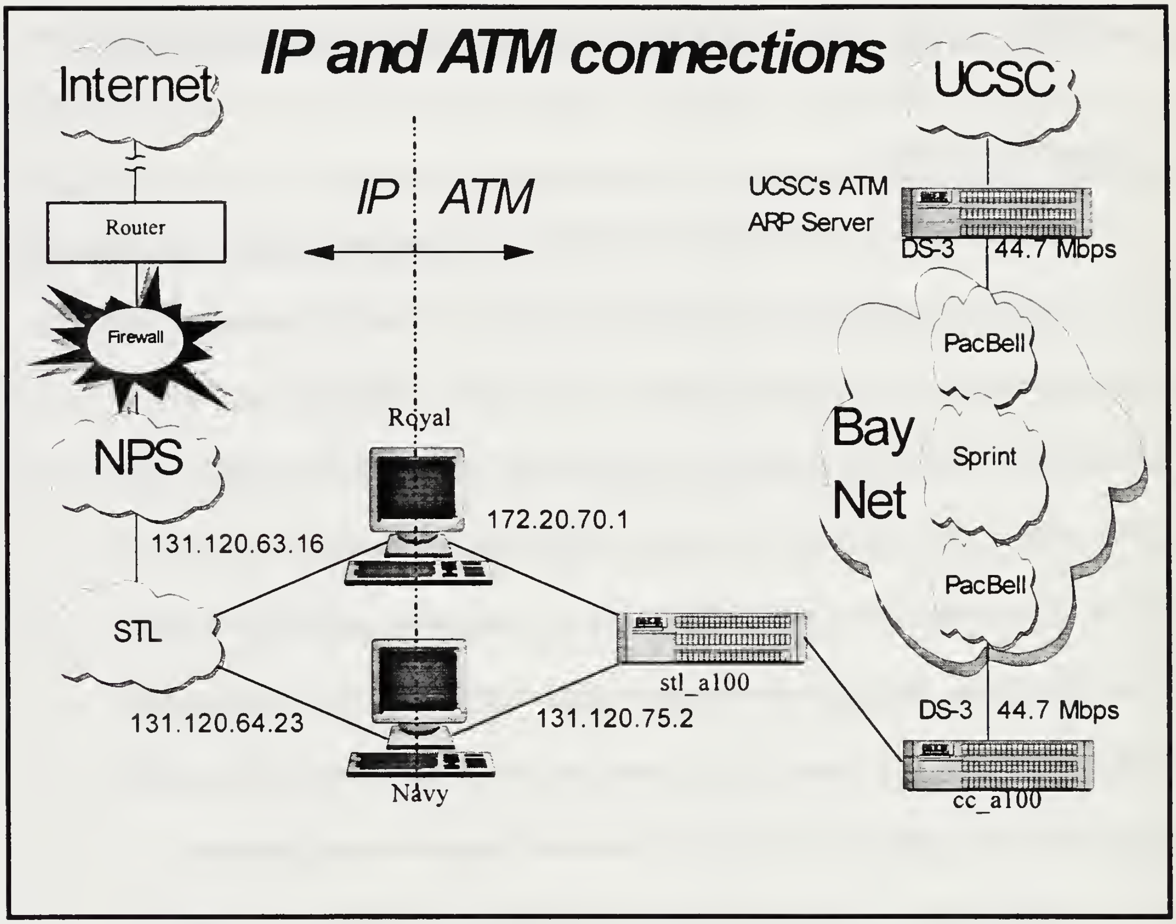

Figure 9. IP/ATM LAN/WAN topology. From (Courtney, 1996).

and then to the Internet through the campus firewall. On the right side of Figure 9, the ATM interfaces are connected to a Cisco A-100 ATM switch in the STL which is connected to another A-100 ATM switch located in the campus main computer center. From the computer center our ATM link connects directly to the Monterey BayNet regional ATM testbed.

C. SECURITY ASSESSMENT

The observations and assessments made in this section are based on the issues discussed in Chapter IV. Subsection 1 discusses administration issues. Subsection 2 discusses environmental issues. Subsection 3 discusses software and data issues. Subsection 4 discusses telecommunication and system access issues. Subsection 5 discusses implementation issues. Interviews include (Cochran, 1996), (McGregor, 1996) and (Williams, 1996).

It is important to note that STL administration, operations and capabilities are among the most challenging on campus (and certainly more challenging than commercial organizations). Staff manning is less than desired, also a common occurrence on campus. This and other anecdotal evidence leads us to believe that STL issues are representative of conditions throughout NPS.

1. Administration Issues

a. System Availability Requirements

The fundamental mission at NPS is graduate student education and research. LAN/WAN connectivity and the availability of unclassified computing resources are essential to that mission, yet system availability requirements for the STL resources are not documented. Assumptions amongst the STL computing staff are informal but consistent. One staff member stated the system availability requirement was 98% (or system down time not to exceed 28 minutes per day), while another was uncertain but stated that system down time was not to exceed 10 minutes a day which equates to a 99.99% availability requirement.

System availability is not currently tracked or explicitly monitored. Because students or faculty are quick to inform staff members of network problems, monitoring and tracking system availability is considered a waste of manpower.

However, monitoring and tracking system availability not only serves to notify system administrators of network problems but also to document the performance and effectiveness of various controls. Automated network monitoring scripts such as those described in (Edwards, 1996) and (Erdogan, 1996) can simplify this task thereby freeing up staff members for other duties.

Without explicit system availability requirements and accurate system performance monitoring, there is no tangible means to evaluate the effectiveness of networked computer system security. The STL staff must explicitly define and document the system availability requirements for the unclassified computing resources. Extending the capabilities of the (Edwards, 1996) and (Erdogan, 1996) automatic monitoring systems to include system availability and analysis results is a worthwhile area for future work.

b. Security and Acceptable Use Policies (AUPs)

Security and acceptable use policies (AUPs) for NPS in general are described in (NPS, 1992) and (NPS, 1995) respectively. No additional policies are specified for the STL.

Limited NPS security policies are posted in the STL but are not enforced. For example, disallowing food and drink around the computer equipment may be a prudent policy, but it is not enforced and therefore is unheeded by STL users who apparently do not see the utility of the policy. In this case nonenforcement correctly implies irrelevance. Nonenforcement simply reinforces a perception that the NPS administration, the STL computing staff, or both, have a myopic view of the scope of computer security or do not understand the need for it with regards to unclassified systems.

The emphasis here is not so much on the enforcement of the limited policies but on their enforceability, adequacy and appropriateness. Enforcement ability is not so much having staff police users actions (a counterproductive approach) as it is having policies that make sense to users as common practice. Policies are not static; they must reflect the nature, culture and environment of the organization. Therefore the STL computing staff ought to critically examine the existing policies. It is worthwhile to identify and eliminate policies that are unimportant, irrelevant or unenforceable. It is worthwhile to have policies which are truly necessary and which reflect the climate and culture of the organization and the needs of all stakeholders. To improve acceptance of the policies, users must be involved in the policy process.

c. Contingency Planning

As discussed in Chapter IV, contingency planning is dependent upon system availability requirements. NPS directives state that a "contingency plan will be developed for each AIS where a disruption of services would have a critical impact on mission accomplishment" (NPS, 1992, pg. 17). Assuming that LAN/WAN connectivity and the availability of unclassified computing resources are required to accomplish the NPS mission, then according to this directive, a contingency plan is required for all unclassified systems.

The STL staff have conducted a rudimentary risk analysis and developed limited contingency plans for their classified resources but have not done so for their unclassified systems. It is safe to assume that availability of the unclassified systems is much more critical than that of the classified systems as the unclassified resources are used more often by more people. Therefore a thorough risk analysis and contingency plan needs to be developed for the unclassified IP-based resources as well. ATM resources are experimental and not mission critical.

d. Roles and Responsibilities

With only one temporary and three permanent members, the STL computing staff is seriously undermanned to adequately accomplish necessary security functions in addition to their other duties and responsibilities. However, the STL network manager has taken steps to obtain external funding and increase the number of staff members to six permanent positions plus the one temporary position currently filled. If approved by NPS administration, this will greatly improve the distribution of labor. Assigning one of these positions the primary responsibility of network security will also help improve the security posture of the STL, since more time might be devoted to network security. To their credit however, the STL computing staff (even though undermanned) does provide exceptional support to the faculty and students.

2. Environmental Issues

a. Natural Disasters

The only environmental threat the STL staff seems even remotely concerned about is excessive heat within the computer room. This is understandable as there already have been some heat-related hardware failures. Steps already have been taken to fund, procure and install an air conditioning system in the STL to remedy this situation. The system is expected to be installed by the end of the year.

While excessive heat is the only concern of the STL staff, other environmental threats certainly exist. The threat of fire is always present. The computer room is carpeted presenting the threats associated with static electricity and dust. Electrical power is not necessarily stable. However, few controls are in place to protect against these other threats. Automatic heat or smoke detectors are required by local directive (NPS, 1992) but are not present. Automatic sprinklers are not installed and the nearest fire extinguisher is outside the computer room and thirty yards down the passageway. Uninterruptible power supplies and surge protectors are not installed.

The risks associated with these other threats may already have been considered but nothing has been documented. Without a formal risk assessment and documented contingency plan, it is impossible to adequately assess the importance of the other threats or obtain the necessary resources to adequately protect against them. As mentioned above, the STL staff must complete a formal risk assessment and develop a contingency plan for their unclassified systems. It is possible that some of these physical safeguards might be funded by campus resources rather than reimbursable research.

b. Human Disasters

The STL computing staff considers authorized system users to be the greatest threat to the security of the STL systems. Ordinarily the typical user is a student or faculty member interested only in their relevant application programs. They give little thought to issues such as password security. Most have only a general knowledge of file permissions and how to effectively use them, despite the requirement for annual computer security training for all personnel. The annual training provided by the NPS administration (typically a short briefing lasting less than 30 minutes) is woefully inadequate. The STL staff cannot rely on this limited training but must take the necessary steps to ensure users are adequately trained.

Physical security is minimal. Laboratory spaces are protected by cipher lock when unmanned. Inside the labs, network hardware is freely accessible. Network cabling and concentrators are not physically secure and hence are vulnerable to either deliberate attack or accidental disconnection by users or cleaning personnel. The main network file servers and ATM switches are also freely accessible. While it is understandable that the ATM switches are accessible due to ongoing research in ATM network technology, it is unwise to allow free physical access to the main network file server. Likewise, the Kerberos key distribution center for the STL test implementation of the Kerberos authentication and authorization system is freely accessible to anyone. This is unsatisfactory for a meaningful implementation of Kerberos. These important pieces of the network architecture need to be physically secure inside ordinarily locked rooms.

3. Software and Data Issues

a. Configuration Management

The STL staff tracks system software configuration reasonably well, but they do not closely track application software installed for course work or research projects. There is no mechanism besides notification by students or faculty that application software is no longer needed. Unneeded application software not only takes up disk space that may be required for other programs or files, but it may also interfere with the proper configuration of newly installed software. This configuration problem will always exist, but removing unnecessary applications may reduce the possibility of misconfiguration to some degree. An administrative policy and an adequate number of administrators are needed.

Another vulnerability exists with the availability of CGI and *perl* scripts on home pages. CGI and *perl* scripts may leak information about the host system configuration that hackers can exploit. The simple fact that a script's source is readable simply by selecting the "View Source" option on the browser allows an intruder the opportunity to study a script and search for ways it can be exploited. Additionally, scripts can inadvertently pass remote user input that contains shell metacharacters to system procedure calls such as *eval()*, *exec()*, *popen()*, *system()*. The point here is not that these system procedure calls should not be use but that script writers need to sanitize user input for proper form and content before passing the input to the shell for execution. With *exec()* and *system()*, script writers can also execute external programs by sending external data directly to the program rather than going through a shell. (Stein, 1996) provides an excellent discussion of the vulnerabilities associated with CGI and *perl* scripts. (Stein, 1996)

Security of CGI and *perl* scripts is only as good as the design and testing of the scripts. The STL staff has implemented adequate procedures to keep scripts inaccessible to external hosts until they have been tested for security. CGI and *perl* scripts are tested and validated on an isolated internal server before being made available on the

STL web server accessible by Internet users. This practice ensures an acceptable level of security within the STL.

b. Integrity and Confidentiality

Confidentiality and integrity of most system software and data are not major concerns for unclassified systems. Integrity and confidentiality are maintained primarily by controlling access to the network (discussed below).

c. System Backups

System backups are woefully inadequate in STL. Full system backups are done only quarterly; incremental backups are attempted weekly. Even when full system backups are conducted only system software is backed up, user data files are not. The attitude and policy of the STL staff is that users must backup their own data files. This is certainly prudent advice and users should attempt to do so to as great an extent as possible, but user backups are infeasible for large files. When questioned about this policy the STL network manager stated that users need only inform the STL staff and appropriate disk space on a mirrored drive will be provided, yet users routinely are not aware of this. This demonstrates an additional need for a user indoctrination and training program for the STL.

Two recent hard disk failures have further demonstrated the need for more frequent and reliable backups. In one case the failure resulted in lost research documentation that had been stored in group accounts. To little avail, one STL system administrator spent the better part of a week trying to recover the lost files. Many of the lost files were written by users who have long since departed and backups were not available. It is unreasonable to request faculty members keep backup copies of all ongoing or completed research documentation on floppy disks. The second case was only a partial failure affecting application files. The files that were affected were able to be moved to another portion of the disk before any files were permanently lost. However, proper reconfiguration of the files has not been guaranteed.

The STL staff need to improve their backup policy. Given the nature of today's technology, system users reasonably expect adequate backups of all files including user data files. Current backup practices within the STL do not sufficiently assure the availability of computing resources and services. There are plans to implement an automated tape backup system with future upgrades to optical disks. This is a step in the right direction.

4. Telecommunications and System Access Issues

a. Access Control

Open computing and maintaining system availability are both essential for the STL. Access security in the STL is good but the network is still vulnerable. While the current Internet connection must pass through the campus firewall, the ATM connection does not. This recognized risk is currently acceptable because only permanent virtual circuits (PVCs) are configured and the risk is minimal. However, as experience with ATM expands and the use of switched virtual circuits (SVCs) become common, this practice will become unacceptable since the ATM call setup procedure may be vulnerable to password compromise or other threats just as with standard Internet connections. If plans to move forward with ATM are implemented, the STL staff must examine and plan ATM security now rather than as an afterthought.

Internal campus network security has also improved considerably since the December 1995 security incident. *TCP_wrapper*[10] is used on all campus traffic and passwords are required for remote login and *telnet* sessions between campus subnets. If a sniffing program is somehow installed inside the firewall, the network is still vulnerable to common password-sniffing attacks. Although not a requirement, sniffer program detection for all NPS LANs is performed periodically. However, passwords are still sent across the network in the clear for all remote login sessions. This is discussed in the next section.

b. Identification, Authentication and Authorization

Password security remains the greatest threat to network security. Passwords are transmitted in the clear over the network whether that network is the campus backbone or the Internet. Additionally, passwords stored on the system are encrypted but the file is world-readable. Although *CRACK*[11] is run periodically to check the security of Unix passwords, there are no guarantees that a "good" password (i.e. one

[10] *TCP_wrapper* is a utility program that allows system administrators to customize their network access security policy by specifying which network hosts are granted or denied access to specific network services as specified in the *host.allow* and the *host.deny* files respectively.

[11] *CRACK* is a utility program that checks the security of user passwords. The program encrypts common dictionary words and then compares these words against the system's password file. If a match is found, the matched password is flagged as insecure. The dictionary can be customized and expanded to meet the needs of the organization. For example, NPS includes ship names and hull numbers in its *CRACK* dictionary. Information on *CRACK* is available at *http://www.ja.net/newsfiles/janinfo/cert/Muffett/Crack-4.1-readme.txt*

not discovered by *CRACK*) has not been compromised by a password-sniffing attack. Additionally, although passwords are required for remote login sessions originating from any machine outside the STL subnet, the default password for a user's STL account is their main NPS Computer Center password. Any compromise of the main campus password file is likely to compromise the STL password file as well. The STL does not require users to use different passwords than their main campus password nor are users required to periodically change their passwords. This is an understandable local policy since network security monitoring is not coordinated campus-wide. Furthermore, the STL does not implement shadow passwords as recommended by CERT (CERT, 1996c).[12]

c. Nonrepudiation

Nonrepudiation is not an identified user need and is therefore not applicable to the STL.

d. Integrity and Confidentiality

Integrity and confidentiality of passwords during network transmission is far more important than integrity and confidentiality of data during transmission. This is because data can be encrypted, but encryption can be broken if passwords are not secure. Preventative and corrective actions are needed because passwords are not protected during network transmission. The option of data confidentiality is a desirable feature for some users but is only supported by ad hoc usage of routines such as the Unix *crypt* command, a trivial encryption routine at best (Curry, 1990).

[12] A shadow password system, does not have encrypted passwords in the password field of the */etc/passwd* file. The encrypted passwords are stored in a shadow file that is not world-readable. (CERT, 1996c)

e. Network Connectivity

(1) IP-to-ATM Connections. Only preliminary security implications of IP and ATM interaction have been considered by the STL and NPS computing center staff. Admittedly the security implications are not clear due to the immaturity of ATM technology. Nevertheless if ATM implementation and use within the STL is to progress, security must be planned in advance rather than added as an afterthought. This is an important area for future study before operational deployment of ATM at NPS. NPS ATM researchers might also consider membership in the closed ATM Forum, but true security will only emerge from open solutions compatible with IP and endorsed by the IETF. In any case, the major ATM deficiencies identified in (Courtney, 1996) must be addressed prior to large-scale use of ATM is advisable.

(2) Direct-dial Telephone Lines. The STL does not have direct modem access. However, the NPS Computing Center and the NPS Computer Science Department subnets do have modem access. Because the STL subnet is accessible from these subnets via the campus backbone, its security is dependent upon the security of these subnets. The STL subnet may be compromised if one of the other subnets are compromised. The STL staff must take the necessary steps to secure their own subnet independent of the security associated with other subnets around the campus. This might be partly accomplished by requiring passwords uniquely different than those used for other subnets. This can also be accomplished if a one-time password mechanism such as S-Key or secure login mechanism such as Kerberos is used for dial-in access on the other subnets.

(3) Multicast Traffic. Multicasting is essential to the work the IIRG is doing with regards to practical internetworking. Because there are no direct security hooks in current multicast protocols, current practice is to use higher-level encryption for data secrecy and/or integrity. The work being done by the IIRG does not require such strict confidentiality and integrity measures. What is required is assurance against denial of service. Certain MBone tools have been made publicly available via home pages using tools such as *mrinfo* and *mtrace*. Granting public user-friendly access to these tools provides important benefits without opening security vulnerabilities as previously assumed. (Erdogan, 1996) provides a detailed examination of these issues, a technical assessment of software vulnerabilities and a secure working implementation which demonstrates the value of this approach.

5. Implementation Issues

a. Effectiveness and Efficiency

There are no metrics in place by which to measure the effectiveness or efficiency of computer security. System availability is not monitored or tracked. Labor hours expended on security related tasks are not documented, nor are the number or frequency of unauthorized access attempts. Until the STL staff either implements procedures or software tools to automatically monitor such metrics, effectiveness and efficiency will go unmeasured. This is a valuable area for further student research.

b. Life-Cycle Costs

Security has not been planned for (or designed into) the unclassified systems, other than out-of-the-box operating system security. There is no documented risk analysis against which to compare costs. Life-cycle costs cannot be evaluated until a complete risk analysis is accomplished. The same metrics used to measure effectiveness and efficiency can be used to track the costs associated with the presence (or lack) of security controls.

c. Convenience and Acceptance

While it is true that security controls must be convenient and accepted by the users if they are to be effective, it is also true that the lack of security controls must be convenient and accepted by the users. The lack of software and data backups has inconvenienced STL users in the past. Unless the present haphazard backup capabilities are changed, significant damage is likely to occur in the future.

The STL computing staff must actively seek both faculty and student input concerning network security. All system users (faculty, staff and students alike) must have a voice in security discussions and decisions. The computing staff needs to actively seek input, because most users will not volunteer such input but will rather complain amongst themselves that service is unsatisfactory.

D. SUMMARY

This chapter briefly describes the integrated IP/ATM LAN/WAN network topology at NPS. Section C documents in detail the observations and conclusions from a site security survey. The conclusions drawn in Section C are largely based on the issues discussed in Chapter IV.

VII. EXPERIMENTAL RESULTS - KERBEROS INSTALLATION

A. INTRODUCTION

This chapter explains the steps taken to install, configure and test Kerberos for our limited implementation. Section B documents the steps taken to obtain the software. Section C briefly describes the steps taken to install the Kerberos software. Section D describes configuration and testing. Section E briefly describes firewall configuration requirements. Section F discusses the significance of remote applications with respect to network security and the steps taken to install and configure a remote application. Section G documents various configuration and technical support options available. Section H describes the methods we used to distribute Kerberos passwords to both local and remote users.

To more easily examine the feasibility of Kerberos we have limited our test implementation to a total of five network hosts and a portable laptop computer. Because the Kerberos protocol is implemented strictly in software, no special equipment is required. Locally we have installed Kerberos on three Unix workstations and a networked PC running Windows 95. Remotely we have installed Kerberos on a portable laptop at NPS (also running Windows 95) and a single Unix workstation at Old Dominion University (ODU) in Norfolk Virginia. We classify the laptop as a remote machine since laptops are typically used for remote computing while traveling. Figure 10 lists the general steps required to implement Kerberos.

- Install Kerberos software on machines at your site
- Set up a Kerberos Key Distribution Center (KDC)
- Configure Kerberos application servers
- Configure Kerberos application clients
- Add users and their passwords to the KDC
- Inform users about Kerberos
- Turn off "non-kerberized" applications (optional)

Figure 10. Steps required to install Kerberos. After (Eichin and McGregor, 1995).

For our limited implementation we established a single realm with the name STL.NPS.NAVY.MIL to match the unique domain name of our LAN. The realm consists of a single master Kerberos Key Distribution Center (KDC), two application servers and six user clients on the six machines described above.

B. OBTAINING THE KERBEROS SOFTWARE

Most Kerberos source code is freely available from MIT and its authorized distribution agents via anonymous *ftp*. While the source code is free, it incorporates 64-bit encryption and is therefore considered a munition by the U.S. Government. As such it is subject to strict export restrictions and cannot be placed on well known *ftp* servers. MIT freeware information is available at:

http://athena-dist.mit.edu/pub/kerberos

We obtained both Kerberos version 4 and Kerberos version 5 software from Cygnus Network Security (CNS). Details are described below.

1. Version 4

Information and instructions for obtaining a copy of the CNS Kerberos version 4 source code is available at:

http://www.cygnus.com/product/cns/sources.html

Simply print and complete the *Cygnus Network Security CNS (V4) Package Request Form* and fax the completed form to Cygnus at the number printed on the form. Cygnus will e-mail the *ftp* location in a few business days. Due to the export restrictions, Cygnus must have a signed form on file before they will release the *ftp* site information. Windows binaries (compatible with Windows 95, Windows 3.1X and Windows NT) are freely available at the *anonymous ftp* site.

In addition to free source code, Kerberos binaries are also available for a fee (from a different *ftp* server) for those systems listed in Figure 11. The fee is for technical support as described in Section E.

2. Version 5

Due to configuration and testing failures associated with the Kerberos version 4 Windows binaries, we purchased a special one-month support contract from Cygnus. The Kerberos version 5 binaries for Solaris and Windows, and the Windows and "C" source

Hardware Architecture	Operating Systems
Sparc	SunOS 4.1.3, Solaris 2.4
HPPA	HP-UX 9.05, 10.01
RX/6000	AIX 3.2, 4.1
SGI	IRIX 5.3
i386 (and up)	Linux

Figure 11. CNS supported systems (CNS, 1996).

code were included with the contract. Once Cygnus received a signed nondisclosure statement and payment for the support contract they provided the *ftp* site and the appropriate information for obtaining support. The details of the support contract are discussed in Section E.

C. INSTALLATION

This section describes installation procedures for the Kerberos version 4 software. Upgrading to Kerberos Version 5 remains as future work.

1. Unix Machines

"Root" access is required to install the Kerberos software. Detailed installation steps are given in (Eichin and McGregor, 1995) and are not repeated here. The documentation is available on-line at:

http://www.cygnus.com/library/cns/install_toc.html

It is advisable to print a copy of the documentation for use during the installation and configuration process, rather than trying to read the documentation on-line while installing the software. Figure 12 lists the general steps required to install Kerberos on a Unix machine.

The source code is distributed as a zipped archive file and must be unzipped and extracted using the Unix *gunzip* and *tar* utility programs respectively. We compiled the source and then installed the software on both Sun and SGI machines. Locally, we have Kerberos installed on a Sun SPARCserver 1000 running Solaris 2.4, a Sun SPARCstation 20 also running Solaris 2.4, and an SGI Indigo^2 workstation running IRIX 5.3. Remotely

- Download Kerberos archive file from *anonymous ftp* site
- Unzip and extract Kerberos files
- Run *usr/kerberos/configure*
- Verify *krb.conf* and *krb.realms* files are correctly configured

Figure 12. General steps required to install Kerberos on a Unix workstation (Eichin and McGregor, 1995).

we have Kerberos installed on a single SGI Indigo^2 Extreme workstation at Old Dominion University (ODU) in Norfolk Virginia, also running IRIX 5.3.

The installation process is partially automated using a configuration script included with the software. By default CNS Kerberos version 4 files are installed in the */usr/kerberos* directory. Although it is possible to install Kerberos in a different directory it is not recommended as installation and configuration becomes problematic. To install the Kerberos software run the configuration script by typing:

```
% /usr/kerberos/install/configure
```

The configuration script will prompt for the realm name and automatically configure the */usr/kerberos/lib/krb.conf* and */usr/kerberos/lib/krb.realms* configuration files. If the realm name is not the same as the Internet domain name for the site then follow the procedures documented in (Eichin and McGregor, 1995) for manually configuring these two files. Regardless, it is advisable to verify the entries in both these files.

The first two lines of the *krb.realms* file will be similar to:

```
company.org MKTG.COMPANY.ORG
.company.org MKTG.COMPANY.ORG
```

The lower case portion of each entry represents the domain name and the uppercase portion is the Kerberos realm name. For our implementation these two lines are modified to read:

```
stl.nps.navy.mil STL.NPS.NAVY.MIL
.stl.nps.navy.mil STL.NPS.NAVY.MIL
```

Likewise, the first two lines in the *krb.conf* file will be similar to:

```
MKTG.CORP.ORG
MKTG.CORP.ORG kerberos.corp.org admin server
```

For our implementation the first two lines of the *krb.conf* file are:

```
STL.NPS.NAVY.MIL
STL.NPS.NAVY.MIL kerberos.stl.nps.navy.mil admin server
```

The first line is automatically added by the configuration script to indicate the default realm for this machine. Do not change this line. The second line identifies the KDC for the realm. Manually edit these files if the lines do not accurately represent your configuration.

2. Microsoft Windows Machines

Installing Kerberos on Windows machines is much more straightforward than installing Kerberos on Unix platforms. (Gilmore and McGregor, 1995) provides detailed instructions for installing the Windows binaries and is available on-line at:

http://www.cygnus.com/library/cns/windows_toc.html

Again, it is advisable to print a copy of the documentation for use during the installation process.

The Windows binaries are distributed as a zipped file that must be unzipped and expanded with the DOS *pkunzip* utility program. Copy the zipped file into the *c:\cns* directory before it is unzipped. This is the default installation directory and must first be created, unless installing an upgrade from an earlier version of Kerberos into the already-existing directory. (Gilmore and McGregor, 1995) provides detailed instructions for creating a CNS program group and the necessary program items within the CNS program group. Locally, we have installed Kerberos on a networked Pentium PC and on a portable laptop for use during travel.

D. CONFIGURATION AND TESTING KERBEROS

Detailed configuration procedures are presented in (Eichin and McGregor, 1995) and are not repeated here. The documentation is available on-line at:

http://www.cygnus.com/library/cns/install_toc.html

Once again, printing a copy for use during the configuration process is helpful. Extensive manual pages are also supplied with the software distribution.

1. Key Distribution Center (KDC)

a. Configuration

Configuration and administration of the KDC is involved yet straightforward. We followed the configuration steps presented in (Eichin and McGregor, 1995) without difficulty. Figure 13 lists the general steps required to create, configure and populate the Kerberos Key Distribution Center (KDC). These procedures require "root" (i.e. superuser) access and are partially automated using initialization and configuration programs. These programs simplify the configuration process by prompting for any required information. The single KDC for our realm is configured on the Sun SPARCserver 1000 described above. As specified in the configuration instructions, we

- Run *kdb_init* to create the KDC database
- Run *kdb_edit* to add the first user
- Configure *kadmin* ACLs
- Populate the KDC database using *kadmin* (or *kdb_edit*)
- Start *kerberos* and *kadmind* daemons
- Update network */etc/rc* file (or equivalent)
- Test for proper KDC configuration with *kinit*

Figure 13. General steps required to create, configure and populate the KDC database (Eichin and McGregor, 1995).

added the alias *kerberos.stl.nps.navy.mil* for the hostname of our KDC to the system */etc/hosts* file.

The first step in the configuration process is to run the *kdb_init* program to create the KDC password database. The first user (i.e. root as software installer) is entered into the database using the *kdb_edit* program to edit the database directly. Additional users also can be added using *kdb_edit* but it is much easier to populate the database using the *kadmin* program. The *kadmin* program is enabled in the next step.

kdb_edit differs from *kadmin* in that *kdb_edit* can access the database only from the machine on which the KDC resides whereas *kadmin* can access the database from any host within the realm. For this reason, access to the database using *kadmin* is controlled by three access control lists (ACLs) that contain the names of individuals who are authorized to add, modify or get information. The ACLs are stored as the files *admin_acl.add*, *admin_acl.mod* and *admin_acl.get* respectively. These files must be manually configured using a text editor.

With at least one user entered in the database the next step is to start the Kerberos authentication server and the database administration server (i.e. *kadmind*). The *kadmind* daemon enables *kadmin* program. Start these processes by typing:

```
%kerberos &
%kadmind -n &
```

The & causes these processes to run in the background. The -n option allows *kadmin* to fetch the master database key from the master key cache file rather than requiring the database administrator to enter the master key each time *kadmin* is run. (Note: The database administrator must still enter their admin password when executing one of the commands from within the *kadmin* program.) Add these two commands to the network */etc/rc* file (or equivalent) to ensure Kerberos is automatically restarted in the event the networked Kerberos authentication server must be rebooted.

The final steps in configuring the KDC are to populate the database and to test the administration and authentication functions of the KDC. We populated the database with a limited number of principals (i.e. users) using the *kadmin* program. The *add_new_key* command (*ank* for short) within *kadmin* is used to add principals to the database.

b. Testing

Basic testing procedures are provided in (Eichin and McGregor, 1995). We initially tested *kadmin* by changing passwords for the database administrator as described in (Eichin and McGregor, 1995). No difficulties were encountered. We also tested *kadmin* by querying the database with the *get_entry* command (*get* for short) to ensure the database had actually been populated. This is an important step as we discovered when Professor Glen Wheless, our research partner at ODU, was unable to initialize to the KDC. Running *get* to troubleshoot the error indicated he had not been entered into our KDC database. Running *ank* and adding Professor Wheless as an authorized network principal (i.e. user) cleared this problem.

We tested the authentication server portion of the KDC by running the *kinit* command from both local and remote hosts. The proper syntax is:

```
%kinit username
```

Initial tests from local machines returned a "Principal unknown" error indicating that the principal (i.e. username) had not been added to the database or that *kinit* was searching the wrong KDC. However we were able to successfully initialize by using the -r option with *kinit* to specify the realm. The proper syntax for this is:

```
%kinit -r username@realm
```

Upon further investigation, we discovered that this is a common error and is documented in (Eichin and McGregor, 1995). To avoid having to use the -r option, add */usr/kerberos/bin* prior to */usr/bin* in your path variable. We discovered quite by accident that this is required because the Solaris operating system is shipped with a default, different version of *kinit* installed in */usr/bin*. Without redefining the path variable we were using the wrong version of *kinit*. Testing *kinit* from remote client machines is discussed below.

2. Application Servers

a. Configuration

For our implementation the Sun SPARCserver 1000 and the Sun SPARCstation 20 are configured as application servers. Again we followed the configuration steps presented in (Eichin and McGregor, 1995) without difficulty. We configured the application servers to support the "kerberized" versions of the Berkeley r-command applications (i.e. *rlogin*, *rcp*, and *rsh*), but we did not disable the standard versions of the applications. Users select one version or another by specifying the complete path for the application on the command line when running the application (e.g. */urs/kerberos/bin/rlogin*), or by placing */usr/kerberos/bin* prior to */usr/bin* in the path variable as discussed earlier. No other "kerberized" applications (e.g. *telnet* or *ftp*) are enabled. Figure 14 lists the general steps required to configure a Kerberos application server.

The first step required to configure a machine as a "kerberized" application server is to add a service instance and password for that server to the KDC. (Recall every "kerberized" service/server combination requires a password.) To configure the

- Add service instance to the KDC using *kadmin*.
- Configure the service on the application server with *ksrvutil add*.
- Change the service password with *ksrvutil change*.
- Update the */etc/inetd.conf* file.
- Test the service.
- Turn off "non-kerberized" versions of the service.

Figure 14. General steps required to configure Kerberos application servers (Eichin and McGregor, 1995).

"kerberized" Berkeley r-command suite on the server *servername* first run *kadmin* then at the *admin* prompt enter the following:

```
admin:   ank rcmd.servername
```

You are prompted for a password. At this point the password can be simple as it soon will be changed to something only the server and the KDC know.

To configure the service on the server itself, first log in to the server as "root" then type:

```
%ksrvutil add
```

Again you are prompted for the service's password. Enter the same password used when adding the service to the KDC. Next run the *ksrvutil change* program to automatically change the simple password to something only the KDC and the application server knows. The correct syntax is:

```
%ksrvutil change
```

Finally, update the */etc/inetd.conf* file as specified in (Eichin and McGregor, 1995) and test the application server.

b. Testing

Full operational testing of the application servers cannot be done until clients are configured. However, simply installing Kerberos on the KDC machine and the application server machine makes them client machines as well. Testing the application servers with client machines is discussed below. Initial communications between the KDC and the application server can be tested prior to running *ksrvutil change* by typing:

```
%kinit rcmd.servername
```

Refer to (Eichin and McGregor, 1995) for troubleshooting hints if *kinit* fails. Destroy all tickets when finished testing. The proper syntax to destroy all tickets is simply:

```
% kdestroy
```

3. Clients

a. Unix Machines

(1) Configuration. As mentioned above, Unix machines can be configured as Kerberos clients simply by installing the Kerberos software as described in Section C. Therefore the two application servers are also configured as user clients. Additionally, the SGI Indigo^2 workstation at NPS and the SGI Indigo^2 Extreme workstation at ODU are configured as clients. Figure 15 lists some basic services that can be included in the */etc/services* file on each client machine. Only *kerberos* and *kpasswd* are necessary to allow principals to initialize to the KDC. *klogin* or *eklogin* is required to enable the "kerberized" *rlogin* application (with or without encryption respectively). Other services or "kerberized" applications (such as *kpop*, a "kerberized e-mail server daemon) can be added as required.

Service	Port/ Protocol	Comments
kerberos	750/udp	# Kerberos (server) udp
kerberos	750/tcp	# Kerberos (server) tcp
kpasswd	761/tcp	# Kerberos "passwd"
klogin	543/tcp	# Kerberos authenticated
rlogin		
eklogin	2105/tcp	# Kerberos encrypted rlogin
kshell	544/tcp	# and remote shell

Figure 15. Kerberos version 4 services and port numbers (Cygnus, 1996).

(2) Testing. Operational testing consists of attempting to access the services provided by the application servers. Prior to testing, ensure */usr/kerberos/bin* is placed in the path variable ahead of */usr/bin*. In our case only the "kerberized" r-commands are configured. Our tests consisted of initializing to the KDC with *kinit* and testing the *rlogin* service from all clients to both application servers. We ran both unencrypted and encrypted *rlogin* sessions. The proper syntax for this sequence of tests is:

```
%kinit username
%rlogin servername          # unencrypted session
%rlogin -x servername       # encrypted session
```

Tests between local clients and the servers proved successful. However, testing the remote client at ODU has been unsuccessful. All attempts to initialize to the KDC up to this point have resulted in a *send_to_kdc: retry count exceeded* error. According (Eichin and McGregor, 1995) this indicates that *kinit* cannot reach the KDC and the most likely cause is a misconfigured *krb.conf*. The *krb.conf* and *krb.realms* files were correct, but to verify these results we deleted the *krb.conf* file and re-ran the *configure* installation and configuration script. The same results were obtained on subsequent tests.

We are fairly certain that the problem is not with the client software itself or a problem with the client configuration because the initialization attempts are being recorded in the *kerberos.log* log file indicating the client is reaching the KDC but the KDC is not responding to the request. Requesting technical support from Cygnus to solve this problem proved to be of little value since our support contract covers only version 5. Time and schedule constraints did not permit resolution of the problem without technical support, and as we move to implement Kerberos version 5 this problem with version 4 remains unresolved.

b. Microsoft Windows Clients

(1) Configuration. (Gilmore and McGregor, 1995) describe the steps required to configure Kerberos for Windows. The *krb.rea* and *krb.con* (DOS versions of the *krb.realms* and *krb.conf* files) must be configured during installation exactly as on the Unix machines. We installed and configured our Windows clients according to (Gilmore and McGregor, 1995). The networked PC and the portable laptop are configured as clients.

(2) Testing. Testing the Windows 95 clients also proved to be problematic. All attempts to initialize to the KDC using a Windows 95 client resulted in the same *send_to_kdc: retry count exceeded* error. Both the *krb.con* and *krb.rea* files (i.e. the DOS equivalents to the *krb.conf* and *krb.realms* files) were configured according to (Gilmore and McGregor, 1995). Posting our problem to the comp.protocols.kerberos USENET newsgroup did not yield any useful responses and technical assistance was not available since CNS advised that Kerberos version 4 was not supported. As with our remote Unix client, this version 4 problem remains unresolved. At the time of this writing we are replacing our Kerberos version 4 Windows implementation and progressing with version 5 implementation.

E. FIREWALL CONFIGURATION

Figure 15 lists various Kerberos services together with the corresponding transport layer protocol (TCP or UDP) and network port numbers that each uses. If your network security architecture includes a firewall, it must be configured to allow these service/port combinations to pass. The IP address for each “kerberized” host may also be required depending on your firewall. The firewall configuration for our implementation is listed in Figure 16. Local clients need not be included in the firewall configuration table. The benefit of Kerberos remains that a particular host does not need to be trusted as the protocol authenticates individual users rather than rely on the trustworthiness of any particular host.

Local Servers: spot.stl.nps.navy.mil 131.120.64.4
azure.stl.nps.navy.mil
131.120.64.5

Remote Clients: garcia.ccpo.odu.edu 128.82.38.141

Service	**Port/ Protocol**	**Comments**
kerberos	750/udp	# Kerberos (server) udp
kerberos	750/tcp	# Kerberos (server) tcp
kpasswd	761/tcp	# Kerberos "passwd"
klogin	543/tcp	# Kerberos authenticated rlogin
eklogin	2105/tcp	# Kerberos encrypted rlogin
kpop	1109/tcp	# Pop with Kerberos
kshell	544/tcp	# and remote shell

Figure 16. Firewall configuration for the NPS Kerberos implementation.

F. REMOTE APPLICATIONS

Successful implementation of Kerberos is measured by the ability of remote users to securely access our system and conduct useful work. Our measure of success was secure access to our system by Professor Wheless and the successful execution on our system of the *Vis5D* application software used by Professor Wheless at ODU.

1. *Vis5D* Download and Installation

The latest version of the *Vis5D* application software (Hibbard, 1996) is available at:

http://www.ssec.wisc.edu/~billh/vis5d.html

Both source and executable code is available. Downloading unknown or unfamiliar executable code from the Internet is not recommended as a general rule, but we chose to download the executables for two reasons. First, inspecting source code for security vulnerabilities or mischievous code is not effective. Second, this particular software is written by a reliable source and was recommended by a trusted colleague.

2. Execution

We tested the *Vis5D* software locally to ensure we correctly installed and configured the software. However, Professor Wheless could not securely access our system to remotely execute the software. Therefore we did not accomplish our remote application goals.

G. CONFIGURATION SUPPORT

While Cygnus distributes the Kerberos source code at no charge, technical support requires payment. Cygnus offers a variety of standard support packages for their Kerberos distribution. These support packages for a single realm are listed in Figure 17.

Due to configuration difficulties with the Kerberos version 4 Windows binaries we contracted for a special one-month support agreement with Cygnus. Because our Kerberos implementation is only experimental at this stage, Cygnus agreed to customize a support contract to fit our particular needs. We purchased a one-month contract for $1,000. Our sole purpose for the support agreement was to obtain assistance with the configuration of the Kerberos version 4 binaries for Windows. We did not discover the similar problem with the remote Unix client at ODU until after the support contract was obtained.

Cygnus relies exclusively on electronic mail to report software problems/bugs and to obtain assistance. An automated report generation program (*send-pr*) is included in the supported distribution of the Kerberos software which is not included in the freeware distribution. This automatic reporting procedure proved to be one more obstacle to hurdle. Due to the short duration of our contract we did not take the time to install *send-pr* but instead used standard electronic mail to obtain assistance. While Cygnus attempts to answer all reports within 24 hours, it was two days before we received a reply to our report.

As mentioned above, our sole purpose for obtaining support was to solve the perceived Kerberos version 4 configuration problem with the Windows 95 clients.

Available Cygnus Kerberos Support Contracts

Single Realm Support:

	1000 Principals	**5000 Principals**	**Enterprise**
University	$ 9,000	$ 15,000	custom
Government	$ 24,000	$ 49,995	custom
Commercial	$ 35,000	$ 60,000	custom

Client Support (cost per client):

	Same as Server	**PC or Mac Client**	**Unix**
University	included	$ 1,500	$ 1,500
Government	included	$ 2,500	$ 2,000
Commercial	included	$ 5,000	$ 2,500

Feature Support (in addition to base prices above):

CNS Developer's Kit (CDK) (Support for Kerberos libraries)	
Existing CDK customer	$ 5,000
Non-CDK customer	$ 15,000
Hand-held Authenticator	$ 5,000
AFS Support	$ 25,000
New dot on Matrix	custom
Multiple realms or slave KDCs at multiple sites	custom

Figure 17. Standard support agreements available from Cygnus for CNS Kerberos distribution (Powers, 1996).

However Cygnus technical support informed us in their reply to our problem report that our contract was for version 5 only. It was our understanding that the support contract was for Kerberos version 4. Cygnus was very helpful and flexible in customizing a contract for us so we can only attribute this misunderstanding to mis-communication between the parties. Therefore while we pursue a version 5 installation our version 4 implementation difficulties remain unresolved. We are removing the version 4 distribution to eliminate the possibility of mismatched software versions.

H. KERBEROS PASSWORD DISTRIBUTION

1. Local Users

Passwords are distributed to local users face-to-face. Users must see a system administrator to establish both a local account and to be entered in the Kerberos database. The Kerberos database entry does not have to be the same as the local account user ID and password. However, making the two the same and including the Kerberos initialize procedure in the login script is recommended. This makes the Kerberos initialization procedure transparent to the user.

2. Remote Users

Passwords for remote users can be distributed over the telephone. Our procedure is to telephone the remote user and have a trusted individual who knows the remote user authenticate them either by their voice or by some unique information only the two share. While the PSN is vulnerable to tapping, we assume the risk of compromise to be negligible.

I. SUMMARY

This chapter explains the steps required to obtain, install and configure the Kerberos version 4 software for both Unix machines and Microsoft Windows machines. Installation and configuration test results are also included. Firewall configuration requirements are discussed in Section E. Section F explains the steps taken to install and test an application used by our remote research partner, Professor Glen Wheless, at Old Dominion University. Execution of the application over a secure channel by Professor Wheless is the primary indication of a successful Kerberos implementation. Section G discusses support contracts offered by Cygnus for their Kerberos version 4 distribution. Section G also explains the specifics and motivation for a customized support contract Cygnus provided for NPS. Section H details the steps taken to distribute Kerberos passwords to both local and remote users.

VIII. CONCLUSIONS AND RECOMMENDATIONS

A. INTRODUCTION

This chapter draws conclusions about Kerberos and makes recommendations concerning the security of the integrated IP/ATM LAN/WAN at NPS. Section B presents our conclusions and Section C presents recommendations for additional and future work.

B. RESEARCH CONCLUSIONS

1. STL Site Security

Despite specific items discussed in Chapter VI, site security for the STL is relatively good. The STL network administrators deserve to be commended for their superlative technical knowledge, dedication to meeting user requirements, and ability to succeed despite improper manning. The STL computing staff have installed the latest security patches for system software, *TCP_wrapper* is used for all intra-campus network communications, the staff plans to install an automated backup system, and an air conditioning system is supposed to be installed before the end of the year.

Based on our observation and site security assessment, most security risks for the STL environment are relatively low. However there are still some vulnerabilities and deficiencies. Recommendations are provided below.

2. ATM Security

Chapter II demonstrated that ATM can be as susceptible to attack as any other network technology; security is still a concern when using ATM and fiber optic cables. ATM technology is immature and new vulnerabilities will likely emerge as experience is gained through research such as that conducted by the IIRG at NPS and research testbeds such as the Monterey BayNet regional ATM network. Much research and testing is still required before security implications associated with combining ATM with other networking protocols such as IP are well known. Open communications among the commercial, scientific and educational communities are essential for discovering and correcting possible vulnerabilities. Open communications are essential if ATM is to live up to its much publicized promises. To date, reports by the ATM Forum have been largely inaccessible and progress by the IETF working group has been incomplete.

3. Kerberos Implementation

We assume that host machines on local-area networks behind a firewall are relatively secure even though some firewalls are vulnerable to address-spoofing attacks. We also assume that the network separating the firewalls is insecure. The purpose of our Kerberos implementation then is to connect two secure hosts across an insecure network cloud, in this case the Internet. Kerberos is well suited for this task and can be an effective addition to an organization's network security architecture. Kerberos is effective against both password-sniffing attacks and address-spoofing attacks. Even if a an intruder is able to penetrate a firewall using an address-spoofing attack, a "kerberized" host behind the firewall is secure. Figure 18 illustrates the value Kerberos can add value to the organization's security architecture. However, Kerberos is not widely used and therefore can be difficult to implement. It is not yet a mature "plug-and-play" software, but requires considerable effort to install and configure. Implementing Kerberos on a large scale (such as a campus) requires considerable planning.

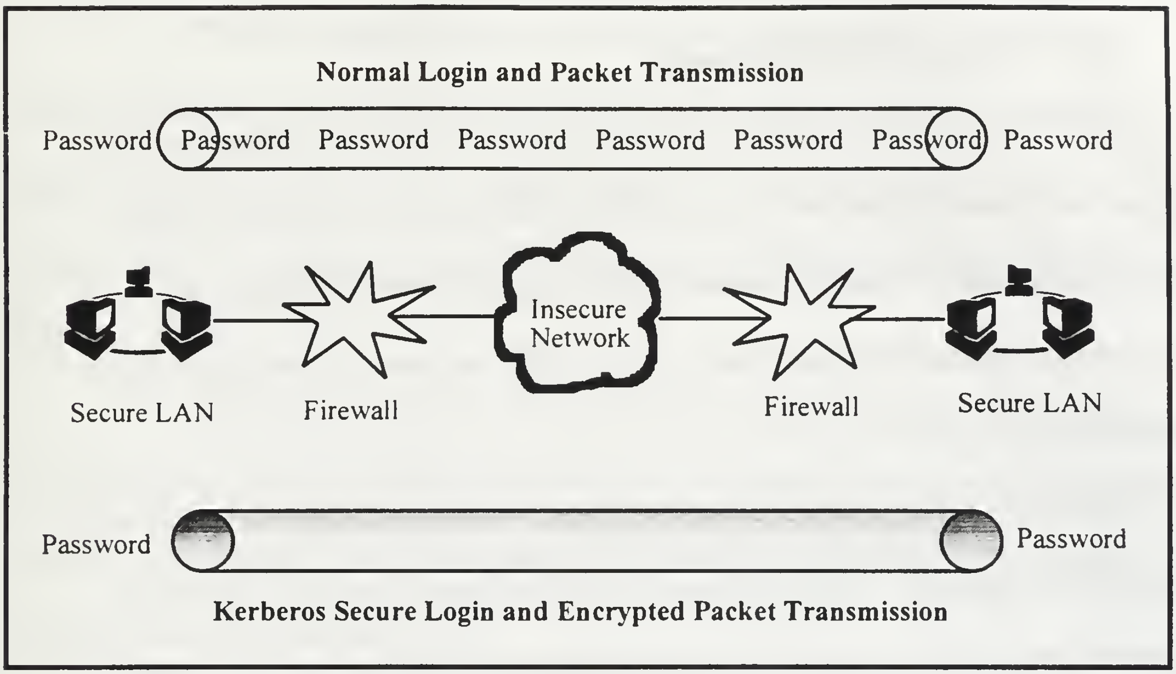

Figure 18. Value of Kerberos within an organization's security architecture.

We limited the scope of this thesis as much as possible from the beginning because (in our experience) implementation projects are seldom as straightforward as anticipated. This proved to be true in this case as well. We have been partially successful with our implementation of Kerberos, but installation and configuration was not as simple as first anticipated. Configuring and operating Kerberos between Unix platforms is relatively straightforward and successful, but configuration of Windows clients was problematic and still has not worked. Our ability to easily operate "kerberized" applications across the Internet between Windows and Unix is required if Kerberos is to be of use to NPS.

C. RECOMMENDATIONS FOR FUTURE WORK

1. Related Work: NPS Theses

The IIRG continues to produce numerous significant theses on internetworking. Thesis security considerations are usually addressed informally. Future IIRG internetworking theses need to include a separate chapter comprehensively addressing security issues. Such an approach has been required for all IETF Requests for Comments (RFC) documents (Postel, 1993).

Restricting access to past or future NPS theses on unclassified network computing security makes little sense. Restricting access to a thesis simply because it documents networked computer system security vulnerabilities removes the motivation to improve security but does not remove the possibility that those vulnerabilities may be exploited. Access to such theses needs to be unrestricted if effective action is to be taken to improve networked computer system security.

2. Site Security

a. Administration Issues

The fundamental security concern for an unclassified system is not necessarily data integrity or confidentiality but rather maintaining system availability. The STL computing staff must determine and document the system availability requirements. They need to conduct a formal risk analysis and prepare detailed contingency plans for the unclassified systems to ensure that system availability requirements are met. Additionally an automatic network monitoring system such as is documented in (Edwards, 1996) and (Erdogan, 1996) needs to be implemented.

The STL computing staff needs to re-examine the current security and acceptable use policies (AUPs) and revise them where necessary to reflect the current environment and culture of NPS. A signed consent form acknowledging the security policies and AUPs ought to be required from all personnel using computing resources at NPS.

It is recommended the STL computing staff implement a user indoctrination program for their unclassified systems similar to the program used for their classified systems. Additionally, it is recommended a first quarter basic computer security course be required for every student at NPS.

The NPS administration needs to expedite funded STL requests for additional staff. Six months between staff departure and opening a relief billet is unacceptably long.

Lastly, documentation concerning the December 1995 computer security breach at NPS is scarce at best. The NPS computer center staff needs to document and widely distribute information concerning the break-in. While this is not to say specific information about the network security architecture at NPS ought to be released, but open communication between computer system administrators, computer users and the NPS administration is essential to avoid similar incidents in the future.

b. Environmental Issues

Physical security within the STL can be improved. The main network file server and the Kerberos authentication/database server are freely accessible to anyone. These two servers need to be moved to a more physically secure environment. Network connections also remain vulnerable. The STL computing staff needs to implement ways to reduce these vulnerabilities.

c. Software and Data Issues

The STL computing staff needs to utilize a configuration management program for application software. The most logical control mechanism is some form of centralized administrative record. At a minimum, the record must include the purpose of the software, the student or faculty member for whom the application is installed, the current version, the current location and the expected review or completion date. Consulting the application log when students depart will allow the STL computing staff to remove any application programs that are no longer required.

The STL network manager needs to accelerate implementation of an automated backup mechanism. With an automated system, complete system backups can be performed monthly with incremental backups performed weekly. Backups must include user data files as well as system files.

d. Telecommunications and System Access Issues

Password security over the Internet remains a significant vulnerability. This vulnerability affects all systems on campus, not just those in the STL. Despite the installation and configuration challenges we experienced, Kerberos is a promising option that should be explored further. NPS also needs to consider implementing a one-time password (OTP) protocol (such as S-Key) for dial-up connections.

3. ATM

If ATM research and implementation at NPS continues, NPS needs to seek membership in the ATM Forum. With an emphasis on practical networking and internetworking, the IIRG and other researchers at NPS can be a powerful voice in the development of interoperable ATM standards.

4. Kerberos

Although the campus network security architecture now includes a firewall and the use of *TCP_wrapper* for all intra-campus network communications, the network is still vulnerable to password-sniffing by undetected sniffer programs and address-spoofing. Kerberos version 4 effectively reduces this vulnerability but it is not without its own weaknesses. Many of these weaknesses have been eliminated in version 5. Therefore upgrading the limited implementation of Kerberos in the STL to version 5 is a priority. Continued work is also required with configuring Windows clients.

A final recommendation is for NPS to investigate the feasibility of a campus-wide implementation of Kerberos. It is safe to assume there will be much resistance to the use of Kerberos. Implementing Kerberos will require change, and change is never an easy nut to crack. Much of this resistance is likely to come from system administrators, many of whom view Kerberos as burdensome and obtrusive. However, configuring and managing a Kerberos site does not have to be any more burdensome than managing a non-Kerberos site. It will only become burdensome if system administrators are expected to manage both independently. Kerberos can be as "unobtrusive" as the standard login sequence; users will never know the difference.

It is clear that careful and detailed planning will be essential for a successful campus-wide implementation of Kerberos. Implementing a commercial version of Kerberos including technical support will be much less problematic than implementing a freeware version of the software. NPS needs to investigate such an option.

D. SUMMARY

It is clear that networked computer security is a multifaceted problem, that encompasses both social and technological issues. This duality was illustrated throughout this thesis. There remains much work and research to be done in the areas of networking and networked computer security. This chapter presents conclusions concerning the security of the integrated IP/ATM LAN/WAN at NPS, ATM security and the Kerberos authentication and authorization protocol. Recommendations for continued and future work are also presented. The most significant result presented here is that secure collaborative networked research by scientists and students distributed around the Internet is achievable. A small amount of additional work and testing is required to make that promise a reality.

APPENDIX A. ON-LINE THESIS DISTRIBUTION

Hypertext markup language (HTML) and PostScript versions of this thesis are available at:

http://www.stl.nps.navy.mil/~iirg/dennis/thesis.html

Comments or questions can be addressed to Don Brutzman at NPS. His e-mail address is:

brutzman@cs.nps.navy.mil

APPENDIX B. ON-LINE KERBEROS DISTRIBUTION

A. Kerberos Documentation and Software Distribution

Kerberos software is available as freeware but its distribution within the United States is controlled because of U.S. export restrictions on cryptographic software. Information about obtaining the Kerberos version 4 or version 5 software can be found at the following URLs.

Cygnus Network Security

Obtaining the Software

http://www.cygnus.com/product/cns/sources.html

Documentation

http://www.cygnus.com/library/cns/install_toc.html	for Unix
http://www.cygnus.com/library/cns/windows_toc.html	for Windows

Note: Only version 4 documentation is available. Version 5 documentation is not yet available on-line.

MIT

Obtaining the Software and Documentation

http://athena-dist.mit.edu/pub/kerberos

B. Technical Support

Technical support is not provided with the distribution. Cygnus does offer technical support for a fee for their distribution of both Kerberos version 4 and version 5.

C. Kerberos Related Papers

An excellent clearing house for Kerberos related papers is at:

http://nii.isi.edu/gost-group/products/kerberos/documentation.html

LIST OF REFERENCES

Alles, A., *ATM Internetworking*, Cisco Systems, Inc., San Jose, CA, 1995. Available at *http://cio.cisco.com/warp/public/614/12.html*

Atkinson, R., *IP Encapsulating Security Payload (ESP)*, RFC 1827, Naval Research Laboratory, Washington, DC, 1995. Available at *ftp://ds.internic.net/rfc/rfc1827.txt*

ATM, *The ATM Forum*, ATM Forum WWW page, ATM Forum, Mountain View, CA, 1996. Available at *http://www.atmforum.com/atmforum/atm_introduction.html*

Baker, D., Manning, S., Meyer, K., and Schaeffer, S., "Addressing Threats in World Wide Web Technology," *OnTheInternet*, May/June 1996, The Internet Society, Reston, VA, 1996, pp. 34-46.

Baskerville, R., *Designing Information Systems Security*, John Wiley & Sons Ltd., New York, NY, 1989.

Bellovin, S., "Security Problems in the TCP/IP Protocol Suite," *Computer Communication Review*, vol. 19 no. 2, April 1989, ACM SIGCOMM, Association for Computing Machinery, New York, NY, 1989, pp. 32-48. Available at *ftp://research.att.com/dist/internet-security/117.ps.Z* or *http://www.acm.org/sigcomm/ccr/archive/1989/apr89/smb.ps*

Bellovin, S. and Merritt, M., *Limitation of the Kerberos Protocol, Computer Communication Review*, vol. 20 no. 4, October 1990, ACM SIGCOMM, Association for Computing Machinery, New York, NY, 1990, pp. 32-48. Available at *ftp://research.att.com/dist/internet-security/kerblimit.usenix.ps* or *http://www.acm.org/sigcomm/ccr/archive/1990/oct90/smb.ps*

Bennington, H., "Report of the Network Security Task Force, November 1990," *Computer Security*, hearing before the Subcommittee on Technology and Competitiveness, Committee on Science, Space, and Technology, U.S. House of Representatives, 102nd Congress, 1st Session, June 27, 1991, U.S. Government Printing Office, Washington, DC, 1991, pp. 51-78.

Blum, D., *Your Company and the Internet*, Rapport Communication, 1995-1996, included in *Conference Proceedings: Internet Expo, Web World, Email World Conference and Exposition, San Jose, CA, February 19-21, 1996*, Digital Consulting, Inc., Andover, MA, 1996.

Bollentin, W., "The Safety Zone," *OnTheInternet*, May/June 1996, The Internet Society, Reston, VA, 1996, p. 6.

Boucher, R., opening statement by Rep. Rick Boucher, *Internet Security*, hearing before the Subcommittee on Science, Committee on Science, Space, and Technology, U.S. House of Representatives, 103rd Congress, 2nd Session, March 22, 1994, U.S. Government Printing Office, Washington, DC, 1994, pp. 1-6.

Bradner, S. and Madnick, A., editors, *IPng: Internet Protocol Next Generation*, Addison-Wesley Publishing Company, Reading, MA, 1996.

Buddenberg, R., *Computer Networking and C3I Systems for Emergency Services*, unpublished manuscript, 1995. Available at:
http://dubhe.cc.nps.navy.mil/~budden/book/table_contents.html

Carpenter, B. and Baker, F., *IAB and IESG Statement on Cryptographic Technology and the Internet*, RFC 1984, Internet Society, Reston, VA, 1996. Available at
ftp://ds.internic.net/rfc/rfc1984.txt

CERT, 1994, *CERT Advisory CA-94.01 - Ongoing Network Monitoring Attacks*, Carnegie Mellon University, Pittsburgh, PA, 1994. Available at
ftp://info.cert.org/pub/cert_advisories/

CERT, 1995, *CERT Advisory CA-95.01 - IP Spoofing Attacks and Hijacked Terminal Conections*, Carnegie Mellon University, Pittsburgh, PA, 1995. Available at
ftp://info.cert.org/pub/cert_advisories/

CERT, 1996a, *CERT Advisory CA-96.05 - Java Implementations Can Allow Connections to an Arbitrary Host*, Carnegie Mellon University, Pittsburgh, PA, 1996. Available at
ftp://info.cert.org/pub/cert_advisories/

CERT, 1996b, *CERT Advisory CA-96.07 - Weakness in Java Bytecode Verifier*, Carnegie Mellon University, Pittsburgh, PA, 1996. Available at
ftp://info.cert.org/pub/cert_advisories/

CERT, 1996c, *Protecting Yourself from Password File Attacks*, Carnegie Mellon University, Pittsburgh, PA, 1996. Available at
ftp://info.cert.org/pub/tech_tips/passwd_file_protection

Chuang, S., "Securing ATM Networks," *Cambridge University ATM Document Collection 4*, Cambridge University, Cambridge, MA, 1995. Available at
ftp://ftp.cl.cam.ac.uk/public/papers/ATM/docs-95-10/12.ps.gz

Claffy, K., Braun, H., Polyzos, G., "Tracking Long-Term Growth *of the NSFNET*," *Communications of the ACM*, vol. 37, no. 8, August 1994, Association for Computing Machinery, New York, NY, 1994, pp. 34-45.

CNO, *Copernicus . . . Forward: C4I for the 21st Century*, Chief of Naval Operations (N6C), Navy Office of Information, Washington, DC, 1995. Available at *http://www.navy.mil/navpalib/policy/coperfwd.txt*

CNS, *Cygnus Product Info: Cygnus Network Security*, Cygnus Network Security WWW page, Cygnus, Inc., 1996. Available at *http://www.cygnus.com/product/cns/platform.html*

Cobb, S., *The Stephen Cobb Complete Book of PC and LAN Security*, Tab Books, Blue Ridge Summit, PA, 1992.

Cochran, M., interview with Milena Cochran (STL System Administrator), May 1996.

Comer, D., *Internetworking with TCP/IP Volume I: Principles, Protocols, and Architecture*, Prentice Hall, Upper Saddle River, NJ, 1991.

Courtney, D., *Internetworking: The Naval Postgraduate School (NPS) ATM Local-Area Network (LAN)*, Master's Thesis, Naval Postgraduate School, Monterey, CA 1996. Available at *http://www.stl.nps.navy.mil/~iirg/courtney/thesis.html*

Curry, D., *Improving the Security of Your Unix System*, SRI International, Menlo Park, CA, 1990. Available at *http://www.sri.ucl.ac.be/SRI/documents/unix-secure*

Cygnus, *Kerb*Net Installation Guide*, Release: 0.9, Document Edition: 0.91 beta, Cygnus Support, Mountain View, CA, 1996.

Deering S. and Hinden, R., *Internet Protocol, Version 6 (IPv6) Specification*, RFC 1883, Xerox PARC and Ipsilon Networks, Palo Alto, CA, 1995. Available at *ftp://ds.internic.net/rfc/rfc1883.txt*

DoD, *Department of Defense Standard: Department of Defense Trusted Computer Systems Evaluation Criteria*, DOD 5200.28 STD, December 1985, U.S. Government Printing Office, Washington, DC, 1985

DoN, *Department of the Navy Automated Information Systems (AIS) Security Guidelines*, published under the authority of SECNAVINST 5239.2, 1991. Available at *http://www.cs.nps.navy.mil/curricula/tracks/security/gs_011B.html*

Edwards, E., *Internetworking: Global and Local Network Monitoring*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996. Available at *http://www.stl.nps.navy.mil/~iirg/edwards/thesis.html*

Eichin, M. and McGregor, P., *Cygnus Network Security: Installation Notes*, Cygnus Support, Mountain View, CA, 1995. Available at *http://www.cygnus.com/library/cns/install_toc.html*

Erdogan, R., *Internetworking: Implementation of Multicast and MBone over Frame Relay Networks*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996. Available at *http://www.stl.nps.navy.mil/~iirg/erdogan/thesis.html*

Fahn, P., *Answers to Frequently Asked Questions About Today's Cryptography*, part 3, RSA Laboratories, Redwood City, CA, 1993. Available at *ftp://ftp.rsa.com/pub/faq*

Fitzgerald, J., *Business Data Communications: Basic Concepts, Security, and Design*, 4th edition, John Wiley & Sons, Inc., New York, NY, 1993.

Flemming, R., *Vulnerability Assessment Using Fuzzy Logic Based Method*, Air Force Institute of Technology, Wright-Patterson AFB, OH, 1993.

Fogleman, R. R. and Widnall, S. E., *Department of the Air Force: Global Presence 1995*, U.S. Government Printing Office, Washington, DC, 1995.

Franklin, J., electronic mail correspondence between Jeff Franklin (NPS AIS Security Officer) and Ronald Dennis, 03 September 1996.

Fraser, B., editor, *Site Security Handbook*, Internet Draft, Carnegie Mellon University, Pittsburgh, PA, June 1996, work in progress expires January 1997. Available at *ftp://ds.internic.net/internet-drafts/draft-ietf-ssh-handbook-03.txt*

Gilmore, J. and McGregor, P., *Cygnus Network Security: Microsoft Windows Client User's Guide*, Cygnus Support, Mountain View, CA, 1995. Available at *http://www.cygnus.com/library/cns/windows_toc.html*

Haller, N., *The S/KEY One-Time Password System*, RFC 1760, Bellcore, Morristown, NJ, 1995. Available at *ftp://ds.internic.net/rfc/rfc1760*.txt

Haller, N., *The S/KEY One Time Password System*, Bellcore, Morristown, NJ, 1994. Available at *ftp://thumper.bellcore.com/pub/nmh/docs/ISOC.symp.ps*

Heyne, P., *The Economic Way of Thinking*, 7th edition, Macmillan College Publishing Company, Inc., 1994.

Hibbard, W., *Vis5D Version 4.2* visualization software, University of Wisconson-Madison Space, Science and Engineering Center, Madison, WI, 1996. Available at *http://www.ssec.wisc.edu/~billh/vis5d.html*

Hughes, J., *Combined DES-CBC, HMAC and Replay Prevention Security Transform,* Internet Draft, June 1996, work in progress expires January 1997. Available at *ftp://ftp.ietf.org/internet-drafts/draft-ietf-ipsec-esp-des-md5-02.txt*

Huitema, C., *IPv6: The New Internet Protocol*, Prentice Hall, Upper Saddle River, NJ, 1996.

IETF, *Welcome to the Internet Engineering Task Force*, IETF WWW page, IETF Secretariat - Corporation for National Research Initiatives, Reston, VA, 1996. Available at *http://www.ietf.org*

IIRG, *Information Infrastructure Research Group* IIRG WWW page, 1996. Available at *http://www.stl.nps.navy.mil/~iirg*

ISO 7498-2-1988(E), *Information Processing Systems - Open Systems Interconnection Reference Model - Part 2 Security Architecture.*

Karin, S., Chair, Supercomputing '95, *Proceedings of the 1995 ACM/IEEE Supercomputing Conference*, Association for Computing Machinery, New York, NY, 1995. Available at *http://www.supercomp.org/sc95/proceedings/*

Karn, P., Metzger, P., and Simpson, W., *The ESP DES-CBC Transform*, RFC 1829, Qualcomm, Inc., San Diego, CA, Piermont Information Systems, Inc., New York, NY, and Daydreamer Computer Systems Consulting Services, Madison Heights, MI, 1995. Available at *ftp://ds.internic.net/rfc/rfc1829.txt*

Kirkpatrick, K., "Modeling A LAN Security Server," *Lecture Notes in Computer Security: Local Area Network Security*, Workshop LANSEC '89 Proceedings, pp. 113-137, European Institute for System Security (E.I.S.S.), 1989.

Klaus, C., *Denial of Service Attack Info*, posting to USENET Newsgroup comp.security.misc by Christopher William Klaus, Internet Security Systems, Inc., Atlanta, GA, 21 May 1996.

Kluepfel, H., "Inside Out You Turn Me," *OnTheInternet*, pp. 19-23, May/June 1996, The Internet Society, Reston, VA, 1996.

Kohl, J. and Neuman, C., *The Kerberos Authentication Service (Version 5)*, RFC 1510, Digital Equipment Corp., Nashua, NH and University of Southern California, Marina del Rey, CA, 1993. Available at *ftp://ds.internic.net/rfc/rfc1510.txt*

Kohl, J., Neuman, C. and Ts'O, T., *The Evolution of the Kerberos Authentication Service*, Massachusets Institute of Technology, Cambridge, MA, 1992. Available at *ftp://athena-dist.mit.edu/pub/kerberos/doc/krb_evol.PS*

Krol, E., *The Whole Internet*, 2nd Edition, O'Reilly & Associates, Sebastopol, CA, 1994.

Kumar, V., *MBone: Interactive Multimedia on the Internet*, New Riders Publishing, Redwood City, CA, 1996.

Lottor, M., *Internet Domain Survey, January 1996*, Network Wizards, 1996. Available at *http://www.nw.com*

Lottor, M., *Internet Growth (1981-1991)*, RFC 1296, SRI International, Menlo Park, CA, 1992. Available at *ftp://ds.internic.net/rfc/rfc1296.txt*

Lucas, J., *Ensuring a C2 Level of Trust and Interoperability in a Networked Windows NT Environment*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996.

Macedonia, M. and D. Brutzman, "MBone Provides Audio and Video Across the Internet," *IEEE Computer*, vol. 27 no. 4, April 1994, pp. 30-36, Institute of Electrical and Electronics Engineers, New York, NY, 1994. Available at *ftp://taurus.cs.nps.navy.mil/pub/mbmg/mbone.html*

McGregor, D., interview with Don McGregor (STL System Programmer), July 1996.

McNulty, L., "Security on the Internet," *Internet Security*, hearing before the Subcommittee on Sience, Committee on Science, Space, and Technology, U. S. House of Representatives, 103rd Congress, 2nd Session, March 22, 1994, pp. 55-74, U. S. Government Printing Office, Washington, DC, 1994.

Miller, S. , Neuman, C., Schiller, J., and Saltzer, J., "Kerberos Authentication and Authorization System," *Project Athena Technical Plan*, Section E.2.1, Massachusetts Institute of Technology, Cambridge, MA, 1987.

Morales, J., *Tactical DMS: A Global Broadcast Service Option*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996. Available at *http://dubhe.cc.nps.navy.mil/~seanet/tacDMS/Title.htm*

NCSC, *A Guide to Procurement of Trusted Systems: Language for RFP Specifications and Statements of Work - An Aid to Procurement Initiators*, 30 June 1993, National Computer Security Center, Fort George G. Meade, MD, 1993.

Neuman, C., and Steiner, J., *Authentication of Unknown Entities on an Insecure Network of Untrusted Workstations*, Massachusetts Institute of Technology, Cambridge, MA, 1988. Available at *ftp://athena-dist.mit.edu/pub/kerberos/doc/unix-security.PS*

Nierle, J., *Internetworking: Technical Strategy for Implementing the Next Generation Internet Protocol (IPv6) in the Marine Corps Tactical Data Network*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996. Available at *http://www.stl.nps.navy.mil/~jenierle/thesis.html*

NIST, 1994a, *Federal Information Processing Standards Publication 186: Digital Signature Standard (DSS)*, National Institute of Standards and Technology (NIST), U.S. Department of Commerce, Springfield, VA, 1994. Available at *http://jcdbs.itsi.disa.mil:5580/A1E4053T2839/fip186.zip*

NIST, 1994b, *Federal Information Processing Standards Publication 190: Guidelines for the Use of Advanced Authentication Technology Alternatives*, National Institute of Standards and Technology (NIST), U.S. Department of Commerce, Springfield, VA, 1994.

NIST, *Federal Information Processing Standards Publication 46-2: Data Encryption Standard (DES)*, National Institute of Standards and Technology (NIST), U.S. Department of Commerce, Springfield, VA, 1988. Available at *http://jcdbs.itsi.disa.mil:5580/A1E2989T2839/fip46-2.zip*

Norman, D., electronic mail correspondence between Dave Norman (Director, W. R. Church Computer Center) and Ronald M. Dennis, 03 September 1996.

NPS, 1995a, *NPS Mission, Vision, and Guiding Principles,* Naval Postgraduate School, Monterey, CA, 1995. Available at *http://www.nps.navy.mil/mission_vision.html*

NPS, 1995b, *Policy on Appropriate Use for Users of NPS Computing and Information Systems*, NAVPGSCOLINST 5230.4, 30 March 1995, Naval Postgraduate School, Monterey, CA, 1995. Access is restricted to internal network hosts. Available at *http://www.nps.navy.mil/internal/policy/app_use_policy.html*

NPS, *Naval Postgraduate School Automated Data Processing Security Program*, NAVPGSCOLINST 5239.1A, 30 November 1992, Naval Postgraduate School, Monterey, CA, 1992.

Palmer, I. and G. Potter, *Computer Security Risk Management*, Van Nostrand Reinhold, New York, NY, 1989.

Partridge, C., *Gigabit Networking*, Addison-Wesley Publishing Company, Reading, MA, 1994.

Peyravian, M. and Herreweghen, E., *ATM Security Scope and Requirements*, ATM Forum Technical Committee Report 95-0579, IBM Corporation, 1995. Available at *http://www.zurich.ibm.com/Technology/Security/extern/ATM/*

Postel, J., *Instructions to RFC Authors*, RFC 1543, Information Sciences Institute, University of Southern California, Marina del Rey, CA, 1993. Available at *ftp://ds.internic.net/rfc/rfc1543.txt* or *ftp://ds.internic.net/rfc/rfc-instructions.txt*

Postel, J., *Transport Control Protocol: DARPA Internet Program Protocol Specification*, RFC 793, Information Sciences Institute, University of Southern California, Marina del Rey, CA, 1981. Available at *ftp://ds.internic.net/std/std7.txt* or *ftp://ds.internic.net/rfc/rfc793.txt*

Postel, J., *User Datagram Protocol*, RFC 768, Information Sciences Institute, University of Southern California, Marina del Rey, CA, 1980. Available at *ftp://ds.internic.net/std/std6.txt* or *ftp://ds.internic.net/rfc/rfc768.txt*

Powers, K., electronic mail correspondence between Kathy Powers (Cygnus Support) and Don Brutzman, 1996.

Putcher, F., Posch, R. and Huy, H. A., *Quality of Services and Security Issues in Distributed Applications over ATM*, Institute for Applied Information Processing and Communications Technology, Graz University of Technology, Graz, Austria, 1996.

Ranum, M., maintainer, *Internet Firewalls Frequently Asked Questions*, Marcus J. Ranum, V-ONE Corporation, 1995. Available at *http://www.v-one.com*

Raynovich, R., “Internetworking Security: Building a Firewall for ATM,” *LANTimes*, October 1995. Available at *http://www.wcmh.com/lantimes/95oct/510c038a.html*

Reeves, C., "Security Requirements, Control Objectives, and the Evaluation Role of the Auditor," *Proceedings of the Sixth Seminar on the DoD Computer Security Initiative*, pp. 87-99, National Bureau of Standards, 1983.

Rendleman, J., “Forum to Bring Security to ATM,” *Communications Week*, July 10, 1995, p. 1.

Rich, L. D., *Unix Security: A Penetration Analysis of Navy Computer Systems*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1992. Access restricted to DoD and DoD contractors.

Rivest, R., *The MD5 Message-Digest Algorithm*, RFC 1321, Massachusetts Institute of Technology, Cambridge, MA, 1992. Available at *ftp://ds.internic.net/rfc/rfc1321.txt*

Rivest, R., Shamir, A. and Adleman, L., "A Method for Obtaining Digital Signatures and Public Key Cryptosystems," *Communications of the ACM*, vol. 21, no. 2, February 1978, Association for Computing Machinery, New York, NY, 1978.

Rowe, Gary J., *EMail and On-Line Services Conference: Putting It All Together*, Rapport Communication, 1995, included in *Conference Proceedings: Internet Expo, Web World, Email World Conference and Exposition, San Jose, CA, February 19-21, 1996*, Digital Consulting, Inc., Andover, MA, 1996.

Russell, D. and G.T. Gangemi, Sr., *Computer Security Basics*, O'Reilly & Associates, Inc., Sebastopol, CA, 1991.

Savetz, K., Randall N. and Lepage, Y., *MBONE: Multicasting Tomorrow's Internet*, IDG Books WorldWide, Inc., Foster City, CA, 1996.

Sheth, A., *Implementation of ATM Security Specification 1.0*, Apurva J. Sheth, 1996. Available at *http://www.eecs.ukans.edu/~asheth/classes/850/proj-plan.html*

SPAWAR PD 51 Security Accreditation Plan Guideline (DRAFT), prepared for Department of the Navy Space and Naval Warfare Systems Command by Booz-Allen & Hamilton, Inc., McLean, VA, 1993. Available on-line at *http://infosec.nosc.mil/projects/guidelines/sap.html*

Stallings, W., *Data and Computer Communications*, 4th edition, MacMillian Publishing Company, New York, NY, 1994.

Stein, L., *The World Wide Web Security FAQ*, version 1.2.4, Whitehead Institute for Biomedical Research, Massachusetts Institute of Technology, Cambridge, MA, 1996. Available at: *http://www.genome.wi.mit.edu/WWW/faqs/www-security-faq.html*

Steiner, J., Neuman, B. C., Schiller, J., *Kerberos: An Authentication Service for Open Systems*, Massachusetts Institute of Technology, Cambridge, MA, 1988. Available at *ftp://athena-dist.mit.edu/pub/kerberos/doc/usenix.PS*

Stevenson, D., N. Hillery and G. Byrd, "Security Communications in ATM Networks," *Communications of the ACM*, vol. 38, no. 2, February 1995, Association for Computing Machinery, New York, NY, 1995.

Stoll, C., *The Cuckoo's Egg: Tracing a Spy Through the Maze of Computer Espionage.* Bantam Doubleday Dell, New York, NY, 1989.

Sullivan, *Army Enterprise Strategy: The Vision*, U.S. Government Printing Office, Washington, DC, 1995.

Tabor, D., *IPv6 and Broadband Services Lesson 19*, New Jersey Institute of Technology, 1995. Available at *http://honer.njit.edu:8000/h2/dtabor/public_html/cis456-html/protected/lesson19/single19.html#OUTL19-27*

Tamer, M., *Internetworking: Multicast and ATM Network Prerequisites for Distance Learning*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996. Available at *http://www.stl.nps.navy.mil/~iirg/tamer/thesis.html*

Tiddy, M., *Internetworking: Economic Digital Storage and Retrieval of Digital Audio and Video for Distance Learning*, Master's Thesis, Naval Postgraduate School, Monterey, CA, 1996. Available at *http://www.stl.nps.navy.mil/~iirg/tiddy/thesis.html*

Toensing, V., "Statement of Victoria Toensing, Deputy Assistant Attorney General, Criminal Division, U.S. Department of Justice," *The Computer Fraud and Abuse Act of 1986*, hearing before the Committee on the Judiciary, U.S. Senate, 99th Congress, 2nd Session, April 16, 1986, pp. 15-30, U.S. Government Printing Office, Washington, DC, 1986.

Varma, A., *Introduction to ATM Networks Lecture Notes*, University of California Santa Cruz, Santa Cruz, CA, 1995.

Walker, S. T., "Testimony by Stephen T. Walker," *Computer Security*, hearing before the Subcommittee on Technology and Competitiveness, Committee on Science, Space, and Technology, U. S. House of Representatives, 102nd Congress, 1st Session, June 27, 1991, pp. 32-40, U. S. Government Printing Office, Washington, DC, 1991.

Williams, T., interview with Terry Williams (STL Network Manager), August 1996.

Winkler, I., *Case Study: Social Engineers Wreak Havoc*, Science Applications International Corporation, Annapolis, MD, 1995.

Wong, K. and Watt, S., *Managing Information Security: A Non-technical Management Guide*, Elsevier Science Publishers Ltd., Oxford, England and Computer Weekly Publications, Surry, England, 1990.

Wood, C., Banks, W., Guarro, S., Garcia, A., Hampel, V., and Sartoria, H., *Computer Security: A Comprehensive Controls Checklist*, John Wiley & Sons, Interscience Publications, New York, NY, 1987.

Zimmerman, P., *The Official PGP User's Guide*, The MIT Press, Cambridge, MA, 1995.

Zuckermann, M., "Post-Cold War hysteria or a national threat?" *USA Today*, Arlington, VA, June 5, 1996, pg. 1A.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center 2
8725 John J. Kingman Road, Ste. 0944
Ft. Belvoir, Virginia 22060-6218

2. Dudley Knox Library 2
Naval Postgraduate School
411 Dyer Rd.
Monterey, California 93943-5101

3. Dr. Dan Boger 1
Chair, C4 Academic Group
Code CC
Naval Postgraduate School
Monterey, California 93943-5101

4. Don Brutzman 4
Code UW/Br
Naval Postgraduate School
Monterey, California 93943-5101

5. Rex Buddenberg 1
Code SM/Bu
Naval Postgraduate School
Monterey, California 93943-5101

6. Dale Courtney 1
Code 05A
Naval Postgraduate School
Monterey, California 93943-5101

7. Ronald Dennis 1
2721 Old Tulalip Rd.
Marysville, Washington 98271

8. Dr. James Eagle 1
Chair, UW Academic Group
Code UW
Naval Postgraduate School
Monterey, California 93943-5101

9. Dr. Richard Elster 1
Provost and Academic Dean
Code 02
Naval Postgraduate School
Monterey, California 93943-5101

10. Dr. James C. Emery 1
Associate Provost for Computer Information Services
Code 05
Naval Postgraduate School
Monterey, California 93943-5101

11. Jeff Franklin 1
ADP Security Officer
Code 54
Naval Postgraduate School
Monterey, California 93943-5101

12. Dr. Reuben Harris 1
Chair, Systems Management Department
Code SM
Naval Postgraduate School
Monterey, California 93943-5101

13. Dr. Debora Hensgen 1
Code CS/Hd
Naval Postgraduate School
Monterey, California 93943-5101

14. Dr. Ted Lewis 1
Chair, Computer Science Department
Code CS
Naval Postgraduate School
Monterey, California 93943-5101

15. Dr. G.M. Lundy 1
Code CS/Ln
Naval Postgraduate School
Monterey, California 93943-5101

16. Don McGregor 1
Code C3
Naval Postgraduate School
Monterey, California 93943-5101

17. Dave Norman, Code 51 1
Director, W.R. Church Computer Center
Naval Postgraduate School
Monterey, California 93943-5101

18. Gary Porter, Code C3 1
Naval Postgraduate School
Monterey, California 93943-5101

19. Dr. Maxine Reneker, Code 013 1
Director, Dudley Knox Library
Naval Postgraduate School
Monterey, California 93943-5101

20. Roger Stemp 1
Code CS/Sp
Naval Postgraduate School
Monterey, California 93943-5101

21. Dr. Murali Tummala 1
Code EC/Tu
Naval Postgraduate School
Monterey, California 93943-5101

22. Terry Williams 1
Code C3
Naval Postgraduate School
Monterey, California 93943-5101

23. Dr. Geoffrey XieCode 1
CS/Xi
Naval Postgraduate School
Monterey, California 93943-5101

24. CAPT George Zolla, USN, Ret. 1
Dudley Knox Library
Naval Postgraduate School
Monterey, California 93943-5101

25. Dr. Michael J. Zyda 1
Code CS/Zk
Naval Postgraduate School
Monterey, California 93943-5101

26. Arul K. Ananthanarayanan 1
University of California Santa Cruz
225 Applied Sciences Bldg
Santa Cruz, California 95064

27. Dr. Larry Atkinson 1
Center for Coastal Physical Oceanography
Old Dominion University
Norfolk, Virginia 23529

28. Dr. Marti Atkinson 1
University of California Santa Cruz
CIS/CE Board Office
225 Applied Sciences
Santa Cruz, California 95064

29. Randall Atkinson 1
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, California 95134-1706

30. Roland Baker 1
Santa Cruz County Office of Education
Media and Technology Services
809 Bay Avenue
Capitola, California 95010

31. CAPT Ross D. Barker 1
Senior Member, Smart Ship Program
Office of the Chief of Naval Operations
10701 Parkridge Blvd
Reston, Virginia 22091

32. Dr. Carl R. Berman, Jr. 1
Office of the Provost
CSU Monterey Bay
100 Campus Center
Seaside, California 93955-8001

33. Bill Brutzman 1
3 South Kingman Road
South Orange, New Jersey 07079

34. Jeff Bryant 1
Monterey Bay Aquarium
886 Cannery Row
Monterey, California 93940

35. CERT Coordination Center 1
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213-3890

36. Erik Chaum 1
Code 2251, Building 1171
Naval Undersea Warfare Center Division Newport
1176 Howell Street
Newport, Rhode Island 02841-1708

37. Marke Clinger 1
FORE Systems, Inc.
174 Thorn Hill Road
Warrendale, Pennsylvania 15086

38. Dr. Tom Defanti 1
Electronic Visualization Laboratory
University of Illinois at Chicago, MC 154
851 S. Morgan St., Room 1120
Chicago, Illinois 60607-7053

39. Dr. Nancy Giberson 1
Santa Cruz County Office of Education
809 Bay Avenue
Capitola, California 95010

40. Bruce Gritton 1
Monterey Bay Aquarium Research Institute (MBARI)
P.O. Box 628
Moss Landing, California 95039-0628

41. R. Peter Haddad, M/S 2L-3 1
Hewlett-Packard Laboratories
1501 Page Mill Road
Palo Alto, California 94304-1126

42. Dr. Harold Hawkins 1
Office of Naval Research
Cognitive & Neural S&T Division
Program Officer
Attn: ONR 341
800 North Quincy Street, Ballston Tower One
Arlington, Virginia 22217-5660

43. Mike Herbst 1
Far West Laboratory
730 Harrison Street
San Francisco, California 94107-1241

44. Birt Johnson 1
Pacific Bell
340 Pajaro Street, Room 131
Salinas, California 93901

45. Tom Karwin 1
U.C. Santa Cruz
P.O. Box 7600
Santa Cruz, California 95061-7600

46. Dr. Marc Lipman 1
Mathematical, Computer, and Information Sciences Division
Office of Naval Research
Ballston Centre Tower One
800 North Quincy Street
Arlington, Virginia 22217-5660

47. Dr.William J. "Bill" Lennon 1
Lawrence Livermore National Laboratory
7000 East Avenue (L-541)
Livermore, California 94550-9900

48. Syd Leung 1
Pacific Bell
2600 Camino Ramon, Room 35306
San Ramon, California 94105

49. Jeff Liening 1
FC4A/TNAC - Validation Support Branch
203 West Losey Street, Room 2000
Scott AFB, Illinois 62225-5238

50. Jack Linthicum 1
Technical Information Specialist
New Technology Development Division
Office of Engineering and Technology
Federal Communications Commission
1919 M Street N.W.
Washington, DC 20554

51. Brian Lloyd 1
Lloyd Internetworking
3461 Robin Lane
Cameron Park, California 95681

52. Dr. Mike Loescher 1
Special Assistant for Information Warfare
Deputy Assistant Secretary of the Navy (DASN)
Washington, DC 20350-1000

53. Dr. Melanie Loots 1
Associate Director, National Center for Supercomputing Applications (NCSA)
4119 Beckman Institute, MC 251
University of Illinois at Urbana-Champaign
605 East Springfield Ave.
Champaign, Illinois 61820

54. Dr. Michael R. Macedonia 1
Fraunhofer Center for Research in Computer Graphics, Inc.
Vice President for Developing Global Work Environments
167 Angell Street
Providence, Rhode Island 02906

55. Rick Maciel 1
Pacific Bell ATM Project Manager
2600 Camino Ramon, Room 4S155
San Ramon, California 94583

56. Lora Lee Martin, Director 1
Director of Program and Policy Development
University of California, Santa Cruz
Santa Cruz, California 95064

57. Kam Matray 1
Monterey Bay Technology Education Center,
Monterey Peninsula Unified School District
PO Box 1031
Monterey, California 93942-1031

58. Chris May 1
California State University, Monterey Bay
100 Campus Center
Seaside, California 93955-8001

59. Dr. James H. May 1
California State University, Monterey Bay
100 Campus Center
Seaside, California 93955-8001

60. Michael McCann 1
Monterey Bay Aquarium Research Institute
P.O. Box 628
Moss Landing, California 95039-0628

61. Dr. Ray McClain 1
Moss Landing Marine Laboratories
P.O. Box 450
Moss Landing, California 95039

62. Mike Mellon 1
Monterey County Office of Education
Instructional Resources and Technology
PO Box 80851
Salinas, California 93912-0851

63. Dr. Pat Mantey 1
Chair, Computer Engineering
University of California, Santa Cruz
Santa Cruz, California 95064

64. Captain Mark Moranville 1
SEA03J, Dir Integrated IS Group
Naval Sea Systems Command
2531 Jefferson Davis Hwy.
Arlington, Virginia 22242

65. Michael Newman 1
Newman and Associates
24090 Summit woods Drive
Los Gatos, California 95030

66. RADM(Sel) Kate Paige 1
Commander NSWC
10901 New Hampshire Ave.
Silver Spring, Maryland 20903-5640

67. Richard Pearlman 1
Technical Director, Knowledge Network
2150 Webster, Room 440
Oakland, California 94611

68. Kathy Powers 1
Cygnus Support, Western Region
1937 Landings Dr.
Mountain View, California 94043

69. Capt Jim Quiros 1
FC4A/TNAC - Validation Support Branch
203 West Losey Street, Room 2000
Scott AFB, Illinois 62225-5238

70. LCDR Dean Rich 1
Fleet Information Warfare Center
2555 Amphibious Dr.
Norfolk, Virginia 23521-3225

71. Deborah Richards 1
Industry Consultant, Pacific Bell
2460 North Main Street
Salinas, California 93906

72. Emily Routman 1
Exhibit Developer, San Jose Tech Museum of Innovation
145 West San Carlos Street
San Jose, California 95113

73. Kathy Ruthowski 1
Editor, NetTeach News
13102 Weather Vane Way
Herndon, Virginia 22071-2944

74. Dr. John Schill 1
Program Manager
ISO Planning and C3
Defense Advanced Research Projects Agency
3701 N. Fairfax Drive
Arlington, Virginia 22203-1714

75. Dr. Arthur St.George 1
National Science Foundation
4201 Wilson Boulevard
Arlington, Virginia 22230

76. Dr. Fred Siff 1
Associate Vice Chancellor
Communications and Technology Services
University of California, Santa Cruz
1156 High Street
Santa Cruz, California 95064

77. G. (Joe) Simone 1
Executive Director,
FasTrak Advanced Digital Services
2600 Camino Ramon, Room 4S155
San Ramon, California 94583

78. Diane Siri 1
County Superintendent of Schools
Santa Cruz County Office of Education
809 Bay Avenue
Capitola, California 95010

79. Dr. Larry Smarr 1
Director, National Center for Supercomputing Applications (NCSA)
155 Computing Applications Building, MC 476
University of Illinois at Urbana-Champaign
605 East Springfield
Champaign, Illinois 61820

80. Bradley R. Smith 1
Computer Facilities Manager
University of California, Santa Cruz
145 Applied Sciences Bldg
Santa Cruz, California 95064

81. David Sonderegger 1
Pacific Bell ATM Project Manager
2600 Camino Ramon, Room 4S155
San Ramon, California 94583

82. LCDR Brian Steckler, USN, Ret. 1
25381 Carmel Knolls Drive
Carmel, California 93923

83. David Stihler 1
Monterey County Information Systems
1590 Moffett Street
Salinas, California 93905-3341

84. Chris Taylor 1
Director of Computing and Computer Resources (CCR)
California State University, Monterey Bay
100 Campus Center
Seaside, California 93955-8001

85. Dr. Anujan Varma 1
Associate Professor of Computer Engineering
University of California, Santa Cruz
1156 High Street
Santa Cruz, California 95064

86. Jim Warner 1
Network Engineer, University of California Santa Cruz
C A T S/Network & Telco Services
11 Communications Bldg
Santa Cruz, California 95064

87. David Warren 1
Cabrillo College
6500 Soquel Drive
Aptos, California 95003

88. Steve Watkins 1
University of California Santa Cruz
Science Library
Santa Cruz, California 95064

89. Dr. Steve Webster 1
Director of Education
Monterey Bay Aquarium
886 Cannery Row
Monterey, California 93940

90. Dr. Glen H. Wheless 1
Center for Coastal Physical Oceanography
Old Dominion University
Norfolk, Virginia 23529